Server Consolidation Study Findings Report

Prepared for:

California State Chief Information Officer

and

The Technology Services Board

RFO No. DTS 06-106854

Author:
Intel Solution Services

Publication Date:
June 19, 2007

1. Contents

Table of Contents

List of Figures

List of Tables

2. Executive Summary

State computer systems and the potential for their consolidation have often been a question and consideration for many within the State, including the State Chief Information Officer (CIO), the California Legislature, State agencies, and State departments. Advanced information processing within government is a necessity but has also been a growing cost.

In the Information Technology (IT) industry, discussions about server consolidation and virtualization[1] are common topics with many private sector companies engaging in projects to reduce the number of computers they use. Continuous advances in computer performance have created the opportunity for this type of consolidation and resulted in the average low utilization of many computer systems already deployed.

Server consolidation is primarily motivated by a desire for cost reduction. However, server consolidation directly reduces energy usage. This supports the State's power conservation initiatives and the Governor's Green Building Executive Order S-20-04.

Other strategic advantages that server consolidation can provide for the State include:

- Better security
- Support for Continuity of Government
- Enhancements for the California Public Records Act and e-discovery
- Improved use of facilities and resources
- Common services between departments

This study was commissioned to investigate server consolidation for the Executive Branch of California State Government. State agencies and departments outside the Executive Branch were invited to participate by the State CIO, in order to identify as many opportunities for efficiencies as possible. This study was led by IT professionals from Intel Corporation who have a wide range of experience in data center architectures and their computer systems.

Approach

The methodology used was to collect data, analyze it, and form recommendations. Data collection included gathering an inventory of computer servers within departments, conducting a web-based department survey of CIOs, doing background research, and conducting selected topical discussions and interviews.

After completing the data collection, the team performed an analysis of the data based on best practices for server consolidation and the team's expertise. The analysis led to five key recommendations covered in Chapter 8 - Key Recommendations. The team

---

[1] **Virtualization -** a representation of a real machine using software that provides an operating environment that can run or host a guest operating system. Multiple virtual machines, i.e. servers, can be run in software on one physical machine.

also found additional opportunities for consolidation for departments to consider that are presented and discussed in Chapter 9 - Additional Opportunities.

To support the key recommendations, the team developed a Total Cost of Ownership (TCO) model based on data collected from the State and supplemented with industry data where State data was not readily available. This TCO model was then used for a cost and value analysis for four (4) of the key recommendations.

Growth Rates

Historic growth rates for installed servers within the State government are not known. In addition, the State does not have a projection for future growth rates. The industry average growth rate for the installed base of U.S. volume servers[2] over the past five years has averaged 12.4% annually. The estimated growth rate for the installed base of U.S. volume servers for the next four years is 9.9%[3]

We assume that tight IT budgets will limit the addition of new servers at the State. However, IT budgets will have funds to provide for server refresh. It is our estimation that without a special effort at reduction or consolidation; over the next five years the State will have:

- 5% annual growth rate of installed servers. This is 50% of the estimated 9.9% growth rate of U.S. volume servers and is based on assuming tighter budgets in California Government.
- 15% annual growth rate of online storage.

With current efforts by departments and a statewide focus on server consolidation, the installed base of servers could be reduced. A manageable three-year goal would be:

- 15% reduction of existing servers.
- 20% conversion of physical servers to virtual servers.

Summary of Data – Current Environment

Data Center Facilities

The team performed an on-site walkthrough review of the State's Gold Camp data center and held discussions concerning other facilities and plans to determine capability and capacity for hosting additional State computer systems if needed.

The team believes that with minor rearrangements and small modifications, the Gold Camp data center can play a significant long-term role in the State's server consolidation strategy. Gold Camp is a relatively new state-of-the-art facility with effective raised floor

[2] IDC - Volume server market (consisting of all systems with an average sales value [ASV] below $25,000)
[3] Calculated from IDC data; see Growth Rate chapter in body of report for more information.

utilization below 50%, and power and cooling utilization as low as 26%. This leaves many opportunities to leverage the existing capabilities of this facility. For longer term planning, this facility can be expanded by build-out of another 20,000 ft^2 of raised floor space.

We intended to use survey and inventory data to assess other departmental computer room facilities. The data collected for this purpose was insufficient for this type of analysis. From interviews and conversations, we understand that there is a variety of computer rooms ranging from large data centers to server closets in office buildings.

Summary of Server Inventory

The server inventory collected information for a total of 6,082 servers. This is a large sample and we speculatively estimate that this number represents between 50% and 60% of the total servers.

A few key points from the analysis of the inventory are:

- Of the 6,082 servers, 5,753 were physical servers and 329 were virtual machines. This shows consolidation activity by some departments and that the adoption of server virtualization is already in progress. However, this is a low number of virtual machines and there is significant opportunity here.
- We attempted to categorize servers by a primary function. Of the 6,082 servers, 1,568 were application servers, 910 were File/Print, 653 were database, 430 were directory, 376 were e-mail, 362 were web, 975 were not classified, and the remainder is in other classifications.
- There were 910 file/print servers. This is a large number for this sample and provides a good opportunity for consolidation.
- There were 376 e-mail servers. This is excessive and is considered for consolidation.

Key Recommendations

The team presents five key recommendations titled:

- Data Centers
- In-Department Consolidation
- File Sharing and Document Management
- E-mail
- Virtualization

For each of these recommendations, we discuss details about our findings, the cost and value of the recommendation, and risks associated with adopting the recommendation.

With the exception of Data Centers, we develop scenarios that illustrate potential costs and savings. These are illustrative only but do represent realistic potential. We have remained conservative in our assumptions. However, actual cost and savings are dependent on a number of design considerations and implementation details.

There is overlap between those servers that would be eliminated in the In-Department Consolidation scenario and the File Sharing and E-mail scenarios. For this reason, it cannot be assumed that the cost and savings are additive across all of the key recommendations.

Data Centers

Recommendation:

The Gold Camp data center is significantly underutilized. A second facility is planned for co-processing, business continuity, and to replace the Cannery data center. No significant addition or expansion should be made to existing departmental computer room facilities. Rather, the Gold Camp and Cannery successor data center facilities should be utilized. Based on State-developed minimum requirements, current computer room facilities should be evaluated to determine if they meet the minimum requirements for security and continuity of operations commensurate with the applications and server functions they support. If not, the processing should be moved to the DTS data centers.

The question of data center consolidation and centralization came up frequently and is a controversial subject for some departments. Full discussion of this subject extends beyond this study and the implications of consolidation. There are organizational roles and responsibilities issues that need to be considered and addressed as part of the broader subject.

For this study, we considered the importance of data center utilization in the efficient use of resources. There are a number of computer room facilities within the State departments. These range from very small server room closets that host just a few servers to large, fully equipped data centers.

The cost savings / avoidance to the State are unknown but presumed to be substantial over time. The potential areas for cost savings / avoidance are:

- Using the existing DTS Gold Camp data center facility can improve utilization and efficiency.
    - Increasing utilization at the data center can be done with minimal additional facilities costs.
    - The underutilization at the data center causes the allocation of the facilities, resources, and staff to be spread over a reduced base of servers. If data center utilization was increased, the per server cost for facilities and services would be reduced substantially.
- Consolidating hardware support personnel builds better expertise across a reduced staff.
- Retrofitting existing departmental computer rooms to meet minimum standards would be costly.
- Building new facilities at one or more departments would be costly.

Key risks with centralizing servers include:

- The cost of implementation

- Increased WAN network requirements—relocating servers to a remote location from the users can increase WAN network requirements.
- Application downtime due to relocation
- DTS Rates—rates need to be comparable and reasonable to customers
- Incorrect determination of data center capacity

In-Department Consolidation

Recommendation:

The State CIO should set a goal to eliminate 15% of existing servers through combining workloads and services over the next three years. In addition, a plan should be developed with the cooperation of the departments to meet this goal. A ***simple*** quarterly tracking spreadsheet/system should be set up to record information by department. The system could track the total number of servers, the number of physical servers, the number of virtual servers, the number consolidated during the quarter, and the number virtualized during the quarter. The spreadsheet and progress should be reported quarterly to the State CIO.

While data center consolidation provides the largest cost avoidance and savings, it does not reduce the number of physical servers. The best opportunity for server consolidation remains with the departments. Server consolidation has been an industry initiative and best practice for several years and several State departments have completed some server consolidation or are consolidating servers now.

A cost and value scenario was developed to consider what the five year impact could be if 15% of servers were eliminated. For this "what-if?" scenario, it was assumed there are 9,000 servers and 1,350 were eliminated at the beginning of the first year. For this scenario, the savings would be $54M over five years. More details are covered in Overview of Cost and Value Analysis below.

Other benefits include:

- Reduction of IT staff workload
- Reduction of facilities requirements—this may extend the use of current facilities
- Consolidation savings can be applied to other projects

There are fewer risks with <u>in-department consolidation</u> but they include:

- Difficulty in combining applications or workloads
- Could create a more complex environment
- Applications on the same system image could interfere with each other and cause outages
- Cost of implementation

File Sharing and Document Management

Recommendation:

Near-Term: Where practical, sites with more than two co-located file servers should review utilization and consolidate these servers to two clustered file servers.

Strategic: Evaluate the potential for applying Wide-Area File Systems (WAFS) technology for remote sites. In addition, conduct an analysis of the costs and benefits of implementing a statewide Enterprise Content Management (ECM) service that can provide a more robust and capable document management capability.

There were 910 file and print servers identified in the sample inventory. Consolidating file and print servers has been a top opportunity and best practice for server consolidation at many companies. With an increased need to better manage documents and other content; many companies are deploying more advanced systems for content management. This report describes the concepts of Wide-Area File Systems and Enterprise Content Management and suggests further investigation.

A cost and value scenario was developed to consider what the five-year impact would be for consolidating file servers to a maximum of two servers per site. For this “what-if?” scenario, 882 file servers from 78 sites were considered for consolidation to 127 servers. For this scenario, the savings would be $26M over five years. More details are covered in Overview of Cost and Value Analysis below.

Other benefits include:
- Simplified file services
- Improved document accessibility and search capabilities
- Reduced overall storage requirements

Key risks with consolidating file servers include:
- Increased WAN network requirements—relocating file servers can impact network bandwidth and latency. However, this recommendation advises consolidation within a site.
- Could create a more complex environment
- Cost of implementation

E-mail

Recommendation:

A plan should be developed to convert all departments to a common State e-mail system over a three-year period. Complete the e-mail architecture, engineering, deployment plan, deployment schedule, and pilot in the first year. Convert all departments to the new e-mail system in the second and third years.

In the inventory sample collected, there were 359 servers across 31 departments classified with a primary function of e-mail. Each department is responsible for designing, engineering, deploying, and maintaining their e-mail system. DTS has a cross-department e-mail offering that began in 2005 that a few departments currently use. Alternatively, DTS could evaluate an outsourced e-mail offering for the State if security, confidentiality, service levels, and all requirements can be met.

A cost and value scenario was developed to consider what the five-year impact could be if 250 e-mail servers were consolidated to 30 at the beginning of the first year. For this scenario, the savings would be $11M over five years. More details are covered in Overview of Cost and Value Analysis below.

Other benefits include:

- An engineering team dedicated to e-mail could provide better quality services.
- A higher level of security can be maintained.
- Response to virus and security threats can be quicker without every department having resources responding.
- Standards can be better maintained.
- High availability solutions can be engineered.
- Continuity of Government e-mail strategies would be simpler to engineer and maintain.
- Better support for e-discovery.

Key risks with e-mail consolidation include:

- Network—if not architected properly, this could cause performance issues and increased bandwidth requirements
- Directory—an e-mail directory would have to be synchronized across departments
- Complexity—all State personnel would have to have access to the consolidated e-mail
- The cost of implementation

Virtualization

Recommendation:

The State CIO should set a goal to convert 20% of the existing servers to virtual machines over a three-year period. In addition, a plan should be developed with the cooperation of the departments to meet this goal. A ***simple*** quarterly tracking spreadsheet/system should be set up to record information by department. The spreadsheet and progress should be reported quarterly to the State CIO.

Often, virtualization is the first thought that comes to mind when discussing server consolidation. In fact, it presents the largest opportunity for reducing physical servers. However, virtualization does not eliminate most software licensing costs, operating system costs, and application support costs. We recommend looking at the other opportunities for server consolidation first, as they have the potential for also reducing software licensing, operating system, and application support costs.

A cost and value scenario was developed to consider what the five-year impact could be if 20% of servers were converted from physical machines to virtual machines. For this "what-if?" scenario, it was assumed there are 9,000 servers and 1,800 were converted to virtual machines at the beginning of the first year. For this scenario, the savings would be $14M over five years. More details are covered in Overview of Cost and Value Analysis below.

Other benefits include:

- Backup and recovery could be simplified
- Hardware upgrades could occur without affecting server operating system
- Ability to load-balance
- Provides for simpler DR and Continuity of Government planning and maintenance

There are fewer risks with consolidation through virtualization but they include:

- Over-commitment of physical system resources
- Proliferation of virtual machines
- A more complex environment

Overview of Cost and Value Analysis

The cost and value analysis consisted of TCO analyses for four of the five key recommendations. The total cost of implementing each recommendation is compared to the estimated cost of an approximation of the State's current server environment. The costs are divided into four major expense categories: Hardware, Software, Facilities, and Personnel. Project implementation costs were also included in the analyses. The TCO illustrations demonstrate the potential financial value of implementing each of the key recommendations this study describes.

The server configurations used in the models are proxies to the State's current server environment. We used costs from State pricing sheets when possible and from industry standard sources when appropriate.

Financial Summary

The Key Financial Results from the analyses for each recommendation are:

In-Department - *The estimated total savings from consolidating the workloads of 9000 servers to 7650 servers is $54M—14.3% of the cost to maintain 9000 servers over five years.*

File Server - *The estimated total savings from consolidating file servers from 882 to 127 servers is $26M—72.5% of the cost to maintain 882 file servers over five years.*

E-mail - *The State could potentially save $11M, or 57% over five years of its estimated current e-mail environment costs, by consolidating 250 e-mail servers down to 30 e-mail servers.*

Virtualization - *By virtualizing 1800 servers on 225 new servers the State can potentially save $14M—18.4% over five years of the cost to maintain the 1800 old physical servers.*

Figure 2-1 shows the financial impact the recommendations have on each of the expense categories.

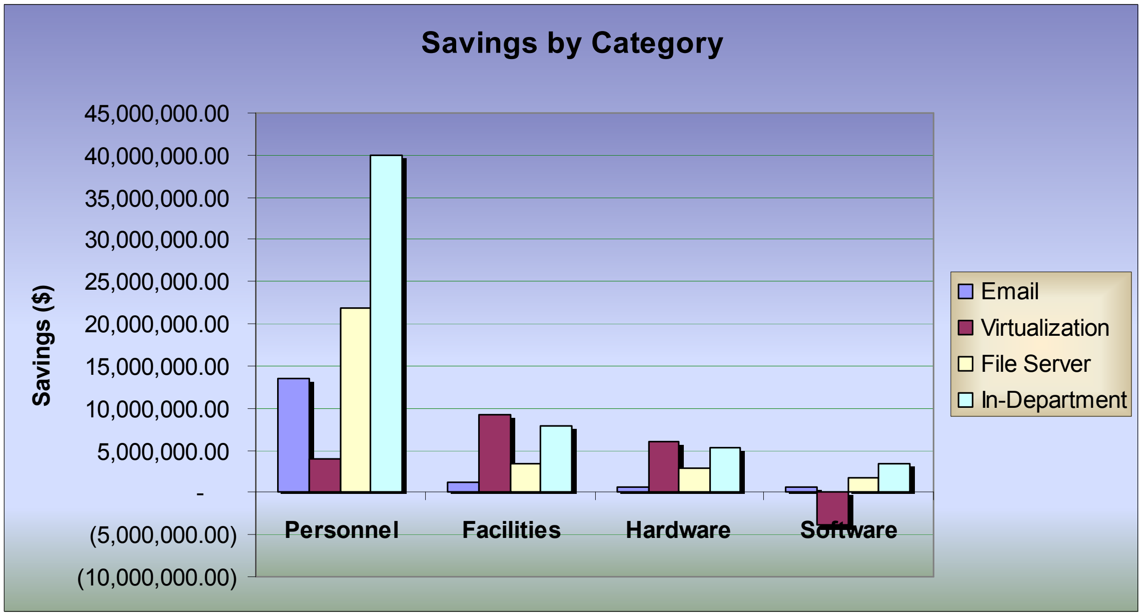

Figure 2-1: Savings per recommendation for each category

The following bullets summarize the impact the recommendations have on each of the four major expense categories:

Software The recommendations affect software costs the least and virtualization actually increases software costs. However, software has the least impact on total cost.

Hardware All of the recommendations will reduce hardware costs, with virtualization realizing the greatest hardware cost savings.

Facilities Savings in facilities costs are generally higher than hardware and software, except in the virtualization model. Facilities costs are reduced by all of the recommendations.

Personnel The greatest impact on costs, and where most recommendations save the most money, is in support personnel; virtualization saves on personnel but not as much as on hardware and facilities. This reduction in demand for personnel allows the State to focus on higher value IT projects and helps address the gap created by State IT personnel retiring or leaving.

Conclusion: Data center consolidation potentially has the largest cost avoidance. All of the other recommendations in this study would save the State money by reducing the number of servers the departments are currently maintaining. According to our financial TCO models, the initiative that would save the State the most money over time is In-Department server consolidation.

Energy Conservation

A welcomed benefit of server consolidation is reduced energy utilization. The State emphasizes and supports energy conservation in its initiatives and programs. Figure 2-2 shows the potential energy savings in megawatt-hours for each recommendation.

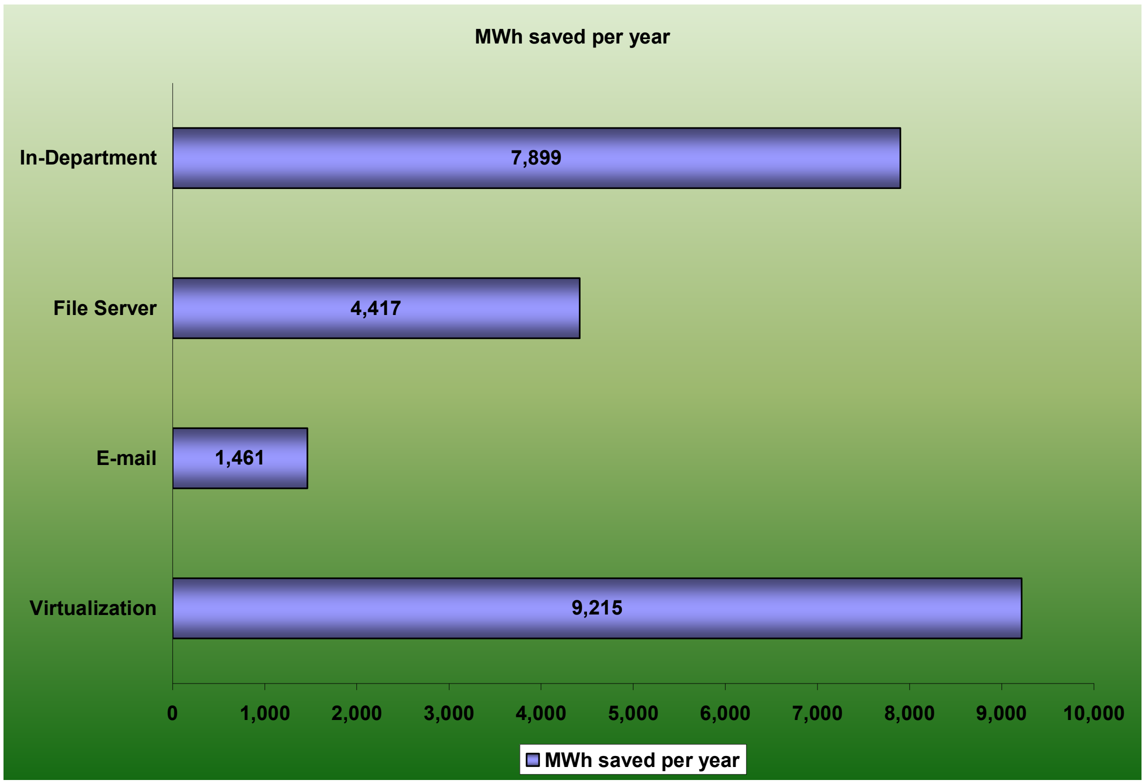

Figure 2-2: Energy savings by recommendation

Issues and Concerns

As the project team performed the data collection, interacted with State staff, and interviewed staff and vendors, several issues and concerns were expressed or noted. These issues and concerns do not represent concerns by all parties or departments but they are significant enough to warrant consideration. We do not provide recommendations on how these issues and concerns should be resolved. However, we agree, they should be addressed as part of any action or project. These departmental/customer issues and concerns are summarized here and further explained in Chapter 10 - Issues and Concerns.

- Centralization
    - DTS cannot provide the same service levels that the departments themselves provide today.
    - DTS rates are higher than what departments can provide for the same service.
    - The overhead of doing business with DTS:

        - adds additional tasks
        - causes undue delays
        - restricts the ability to set priorities
        - slows down projects
- Department Autonomy
    - DTS does not have the same business perspective that the department has.
    - An autonomous IT group can have more control over prioritizing responses to problems according to its own business objectives.
    - An autonomous IT group can monitor and control costs better.
- Network – The network impact of consolidation of servers or services centrally needs to be addressed beforehand.
    - CSGNET backbone capacity — Centralization creates a significant increase in backbone utilization.
    - Network high-availability — If servers or services are moved, then other portions of the WAN topology could become a single-point-of-failure and cause end-user downtime.
    - Endpoint network architecture — Today, end users may not be directly connected to the CSGNET backbone; this could require end-user network changes.
    - Cost — Redesigning WAN connectivity to departments and remote offices, engineering for high-availability, and increases in bandwidth will increase network costs.
- DTS Expertise - to run centralized IT services, DTS must maintain adequate expertise to provide the services.
    - Project Cost / Competing Priorities - there are increased costs, required resources, and staff requirements in the short term that are an impact to departments' budgets.
- Server Funding - many servers were funded specifically by programs, projects, or budget line items. Some of this type of funding is specific and does not provide for consolidated systems.

3. Project Summary

3.1. Overview

The California State Information Technology Strategic Plan, as updated November 2006, calls for the Director of the Department of Technology Services, jointly with the IT Council Technology Services Committee, to submit a Server Consolidation Strategic Plan to the State CIO and the Technology Services Board by June 2007. From the strategic plan:

> Goal 4: "Lower Costs and Improve the Security, Reliability and Performance of the State's IT Infrastructure."
>
> Objective 2: "Consolidate Technology Infrastructure and Service"
>
> "The State will consolidate its technology infrastructure and services to leverage the economies of scale in the utilization of resources, eliminating unnecessary redundancies and reducing support cost through standardization. These efforts will align with the development of the enterprise technology architecture and implement the strategic direction for the use and deployment of information technology solutions statewide. Technology consolidation by the departments and the Department of Technology Services, after consulting with its customers, will increase the security, robustness and reliability of the State's technology infrastructure and improve budget allocation and performance management, cross-agency collaboration, information sharing and e-government solutions."
>
> Action 3:
>
> "By June 2007, the Director of the DTS jointly with the IT Council Technology Services Committee will submit a Server Consolidation Strategic Plan to the State CIO and the Technology Services Board."

This study is in support of developing the Server Consolidation Strategic Plan.

3.2. Scope

The primary scope of this project is a server consolidation study for the Executive Branch of California State Government. State agencies and departments outside the Executive Branch were invited by the State CIO to participate in order to identify as many opportunities for savings and efficiencies as possible.

The objectives of this study are to:

- Create an inventory of existing serves within California state departments (primarily Executive Branch).
- Estimate the future server growth rate
- Identify server consolidation opportunities within targeted agencies of the State of California including quantitative and qualitative improvements to be gained through consolidation and rationalization.
- Determine the cost and value of consolidating servers
- Identify consolidation risks and issues

3.3. Team

This effort was led by a dedicated team of senior engineers and IT professionals at Intel Corporation who have a wide range of experience from engineering “Data Centers of the Future” for leading edge corporations, to solving business and technology problems for governmental organizations.

4. Project Strategy

4.1. Approach

The motivation for a server consolidation project is usually to reduce costs. In addition to cost reductions—Total Cost of Ownership (TCO)—there is an expectation for improvements in Reliability, Availability, and Serviceability (RAS).

Many organizations do expect to both reduce cost AND improve performance through consolidation. With the continued performance improvements of each generation of computers, this is a realistic goal. However, to realize the increasing capacity and performance of new systems, it is necessary to increase the workload applied to these systems.

4.1.1. Cost Reduction

To determine how costs can be reduced, we need to examine the broad categories of costs that factor into the total cost of running applications. These costs include:

Applications and Their Management:

- Application Software Development and/or Licensing – plus configuring, training, maintaining, service agreements, etc.
- Application Support – application administrators, application help desk, monitoring, security, etc.

System Software and Management:

- System Software – OS, utilities, security, monitoring software, backup restore, associated server licensing, etc.
- Systems Support – system administrators, security, patch management, troubleshooting, backup/recovery, etc.

Hardware, Infrastructure, and Facilities:

- System Hardware – dedicated components of a system, including dedicated storage
- Infrastructure – shared components and facilities such as networking, shared storage, floor space, power/thermal, etc.
- Hardware Support – planning, install, de-install, break-fix, maintenance, etc

These are broad categories, and other costs associated with an application, system, or the hardware environment could also be included.

What proportions of costs are associated with each category vary greatly depending on the application, application complexity (number of interfaces and customizations), software licensing costs, degree of automation, and complexity of hardware (high availability and redundancy needs). Usually, support and software licensing is the largest portion of the cost. Hardware cost and support usually contributes a smaller portion of overall TCO.

When prioritizing server consolidation opportunities, we need to give appropriate weight to both cost savings and capability improvements, such as RAS. For example, simple server virtualization does not reduce costs as much as consolidating application instances and removing servers—physical or virtual (as the detailed analysis below will show).

For this study, we do not analyze business application software other than as part of an entire system that can be virtualized.

4.1.2. Consolidation Opportunities

There is a variety of ways to consolidate to reduce costs and improve performance. Listed here are some of the types of consolidation that could be considered:

- Applications:
  - Reduce the number of applications
  - Combine application instances
  - Reduce the number of servers needed for an application instance

  Infrastructure Services:
  - E-mail
  - File and Print
  - Directory
  - Other infrastructure services
- Databases – Combine small and medium database servers
- Web Servers – multiple web sites can be hosted on a single web farm
- Remote Access (Citrix) – multiple terminal services applications can be combined on a Terminal Services farm
- Small and Medium Applications:
  - Shared Landing – combining compatible applications on one system
  - Virtual Machines (VM)
- Development Systems – use virtual machines for development and testing

4.2. Methodology

The following approach for analyzing and developing recommendations was used:

- Data collection
  - Inventory servers
  - Conduct departmental surveys
  - Perform background research
  - Conduct selected topical discussions and interviews
- Analysis
  - Summarize current environment
  - Summarize sites and facilities
  - Determine projected growth rates
  - Create baseline cost summary
  - Determine opportunities from review of Consolidation Opportunities Matrix[4]
- Recommendations

---

[4] Consolidation Opportunities Matrix developed by Intel Solution Services as part of their Server Consolidation Practice.

- Identify key consolidation opportunities
- Identify additional opportunities
- Document issues and concerns

4.2.1. Data Collection

The State of California is responsible for operating a very diverse Information Technology infrastructure that spans the geography of the entire state, and supports all of the activities of state government.

The State's IT assets are controlled by a variety of departments, located in numerous locations. Due to decentralization of servers and management thereof, the project team requested inventories and survey responses from each agency. Data collection was conducted individually for each department by requesting a server inventory and completion of a departmental and site survey for each facility that had servers onsite. In addition, topical discussions and interviews were conducted to gain further background information.

The intent was to obtain participation from 80% of departments expected to have the largest server volume and to obtain inventories of 80% or greater of their servers. Due to a short timeline and departmental resource constraints, we believe we actually captured 50-60% of the actual server inventory through this process. That figure is entirely speculative, since we do not know the total number of servers in use around the State. The captured inventory is large enough to perform analysis and create recommendations.

Since there are many departments, a survey was used to collect information about the IT structure and sites within departments. Due to the short timeline and departmental resource constraints, less than 40% of the departments responded to the online survey. The survey responses provided were used as a sample and assumed to represent the overall environment for the purposes of this study.

The project team met with DTS staff to collect background information on the following topics:

- DTS organization, background, and services
- Facilities and Infrastructure
- Network services
- Directory services and e-mail
- Clients and servers

We also met with selected department-level IT representatives to gather further details concerning department-level IT environments.

4.2.2. Analysis

The project team loaded the server inventory and survey data into a database for analysis. This database resides on a server at DTS and is owned by the State project sponsor.

A summary of the current environment was created based on the sample data and is included in Chapter 7 - Current Environment.

The project team then analyzed the collected data to look for consolidation opportunities. First, the data collected by this study was compared to our standard review list of consolidation opportunities. Next, the project team explored additional consolidation opportunities based upon subjective experience. The team then evaluated these opportunities and created a list of possible consolidation recommendations.

4.2.3. Recommendations

The list of possible consolidation recommendations was further evaluated for impact and feasibility and key recommendations were defined. The key recommendations were then analyzed and a cost and value analysis was completed. Key recommendations are included in Chapter 8 - Key Recommendations.

Opportunities that were not identified as key recommendations are presented and discussed in Chapter 9 - Additional Opportunities. These opportunities should be reviewed by departments and adopted as appropriate.

5. Cost and Value Methodology

The cost and value analysis compares the total cost of ownership (TCO) of the study's recommended solutions to estimates of the current server environments for each scenario. Each scenario's TCO in "Chapter 8 - Key Recommendations" is analyzed separately. The cost of each scenario is an estimate of how much money the State currently spends or, in the case of a recommended solution, what it might spend to implement and maintain the solution. The value is the difference between the costs.

This section describes:

- How the TCO framework is organized
- The cost categories that the framework analyzes
- The source of the financial data
- The costs and assumptions that are consistent across scenarios[5]

5.1. Server Configuration Derivation

The baseline server configurations were derived from our interpretation of the reported inventory data. The recommended server configurations were customized based on the recommendation being analyzed.

5.2. TCO Categories

The TCO models themselves are a standard analysis of total costs including Hardware, Software, Facilities, Personnel, and Implementation costs. The definitions follow:

Table 1: TCO Elements

Hardware	• Cost of **servers** including an additional 2-year warranty beginning the fourth year of ownership (5 years total). • Network costs of one **Gbit network** interface card (NIC) and one host-bus adapter (HBA), used to connect to SAN storage, per server. The servers that are not connected to a SAN will not incur the **HBA** cost. (Note: Servers typically come with an integrated 2-port GbE NIC as part of the base server cost.)
Software	• The **operating system** license is included in the server cost. The operating system assumed for this study was Windows Server 2003 Standard Edition. • Cost of **software utilities** such as virus protection, management agents, and backup. • **Applications** such as SQL Server 2005 for database servers and Exchange Server 2003 for e-mail servers. When possible, the server licensing model was used and CAL licenses were ignored because the number of users

[5] The spreadsheets used to perform this analysis are part of the deliverable for this project. The State is welcome to use those spreadsheets to perform scenario analysis of the impact of different sets of assumptions on the TCO results.

	before consolidation and after consolidation would remain equal.
Facilities	• **Utilities and Power -** costs include the power requirements of each server and the power needed to cool the server infrastructure. Input power is expressed as kilowatt-hours per year and is charged at $.08 per kilowatt-hour. The power requirements of each server were derived from industry standard benchmarks. • **Data Center Construction** – costs include raised floor space, electrical and mechanical. • **Annual maintenance and operations** – yearly costs to maintain and operate the data center. • **Cabling** - costs to pre-wire server racks with copper and fiber. • **LAN/SAN Switch Ports** – costs for Ethernet and fiber channel ports.
Personnel	**Personnel** - costs are based on an Associate Information Systems Analyst (Specialist) State Classification. 36% is added to the base salary to generate a fully burdened rate. This is used as a Full-Time Equivalent (FTE) salary for this study. The financial model assumes that .2 of the FTE is dedicated to hardware support and .8 of the FTE is dedicated to software support. The number of FTEs required to support the server fleet is calculated by multiplying the number of servers by the ratio of FTEs per server.
Implementation	• **Project implementation** – cost includes consultant fees, design costs, and migration costs Specific project implementation costs are highly variable and are subject to numerous factors and design decisions. We supply a judgment cost here to provide a more accurate analysis. • **Server Physical Installations or Removal** – cost of physically racking and installing a new server or of removing and disposing of a server. • **Software Installation** – cost of FTE time to setup and configure the recommended software on new servers.

5.2.1. Hardware Costs and Assumptions

Hardware costs were gathered from Department of General Services' procurement price lists for PC servers. An additional 2-year warranty was added to each server beginning the fourth year of amortization, assuming the first three years were covered by a vendor warranty.

According to the price lists, the cost to add an HBA Fiber Channel card is $617.16.

5.2.2. Software Costs and Assumptions

Department of General Services' procurement price lists do not state the cost of the standard server OS, Windows Server 2003 Standard Edition. Instead, OS costs are included as part of the hardware costs and are itemized as $0 in the TCO models.

Software utilities costs are estimated to be $2,500 per server, spread over five years, and they include costs for virus protection software, management agents, and backup software.

Database and e-mail package costs were gathered from Microsoft's web site and do not include a State-negotiated discount.

Table 2: Software Cost Factors

Application Type	Environment	Application Name	Cost
Operating System	Current	Windows Server 2003 Standard Edition	$999.00 per server
	Recommended	Windows Server 2003 Standard Edition	$999.00 per server
E-mail	Current	Exchange Server 2007 Standard Edition	$699.00
	Recommended	Exchange Server 2007 Enterprise Edition	$3199.20 per server
File/Print	Current	Standard Feature	$0.00
	Recommended	Standard Feature	$0.00
Web		IIS	$0.00
Utilities		Virus protection, management agents, backup	$2,500.00 per server

5.2.3. Facilities Costs and Assumptions

Table 2 shows the assumptions used in the facilities cost calculations of the TCO. Data Center power density was assumed to be 80W per square foot and the cost per kW-hour per year was assumed to be $0.08. The $1150/square foot for data center construction includes raised floor space, and electrical and mechanical installations. The annual maintenance and operating costs are assumed to be $20/square foot. The power for air-conditioning, cooling, and UPS losses is assumed to be the same as the power consumed per server.

Table 3: Facilities Cost Factors

Power and Utilities	
Watts/Square-foot (W/SF)	80
SF/Rack	35

Data Center Construction Costs ($/SF)	1150
Annual Maintenance and Operating Costs ($/SF)	20
KWH (Electricity Unit) Cost ($)	0.08
Power Multiplier to included Air-conditioning/Cooling/UPS Losses	2
Work Days per Year	250
Weekends + Holidays	115
Hours busy	12
Hours Idle	12
Data Center Cabling	
Copper + Fiber / rack - pre-wire	$3,000
48-port GbE line card cost (est. incl. enterprise discounts)	$14,400
FC Port cost (est. incl. Enterprise discounts)	$700
Ethernet ports per server	2
FC Ports per server	1
Data Center Utilization	
Busy Hours per year	3000
Idle Hours per year	5760

5.2.4. Personnel Costs and Assumptions

Personnel costs are based on an Associate Information Systems Analyst (Specialist) State Classification. 36% is added to the base salary to generate a fully burdened rate. This is used as a Full-Time Equivalent (FTE) salary for this study. The financial model assumes that .2 of the FTE is dedicated to hardware support and .8 of the FTE is dedicated to software support.

Note: FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance. FTE's by workload are based on CIOView[6] standard metrics.

Table 4: Personnel costs and assumptions

Generic Server FTE salary	$65,172.00	Per year
Benefits factor	36%	
File/Print Server FTEs	0.066666667 FTE/server	15 servers/FTE
e-mail and comm server FTEs	0.142857143 FTE/server	7 servers/FTE

[6] CIOView Corporation, Acton, Massachusetts

Database server FTEs	1 FTE/server	1 servers/FTE
Web server FTEs	0.066666667 FTE/server	15 servers/FTE
Test and development server FTEs	0.104166667 FTE/server	9.6 servers/FTE
Application	0.1 FTE/server	10 servers/FTE

5.2.5. Implementation

Specific project implementation costs are highly variable and are subject to numerous factors and design decisions. We supply a judgment cost here to provide a more accurate analysis.

The one-time conversion cost is the cost to install or remove physical servers for the recommended solutions plus the time to setup and configure any recommended software solution. The server installation and removal costs were gathered from the Department of General Services' procurement price lists for PC servers.

5.3. TCO Time Period

The time period analyzed in the TCO is five years, reflecting an assumed server refresh rate of five years. Each line item in the TCO model has an amortization period in years over which the cost is divided. The yearly cost is then applied evenly over each year of the analysis. The amortization periods for each category are described below.

- Server hardware costs were amortized over five years. Maintenance costs for hardware that are paid to the vendors were included beginning the fourth year, assuming that the first three years of ownership were covered by vendor warranties.
- Software licensing and maintenance costs were incurred annually.
- Utilities and power costs are paid monthly but are calculated as yearly costs in the TCO.
- The data center construction costs are depreciated over 15 years. Server cabling is depreciated over 10 years.
- LAN/SAN switch ports are depreciated over 5 years.
- Personnel salaries are calculated as yearly costs.
- The conversion cost is a one-time cost paid only the first year of implementation.

Table 5: Amortization Schedule

Hardware	5 yrs
Software	5 yrs
Utilities and Power	1 yr
Data Center Construction	15 yrs
Server Cabling	10 yrs
LAN/SAN Switch ports	5 yrs
Personnel	1 yr
Conversion	1 yr

5.4. Cost and Value Analysis

The cost totals for each year across all categories were aggregated to calculate the TCO for the current and recommended states. The total savings is the difference between the current and recommended TCO models. The value to the State is the total savings gained from implementing the recommended solution.

6. Growth Rates

6.1. Demand for IT Services

The growth rate for IT services in departments is bound by department and program budgets and funding. A constraint on department budgets puts pressure on IT to provide more support services with no additional funding for IT, and sometimes with a reduction in funding for IT. Improving or adding IT services is a way to maintain or reduce overall department costs; however, funding is limited for undertaking new projects.

Figure 6-1 is a graphic representation of IT demand drivers.

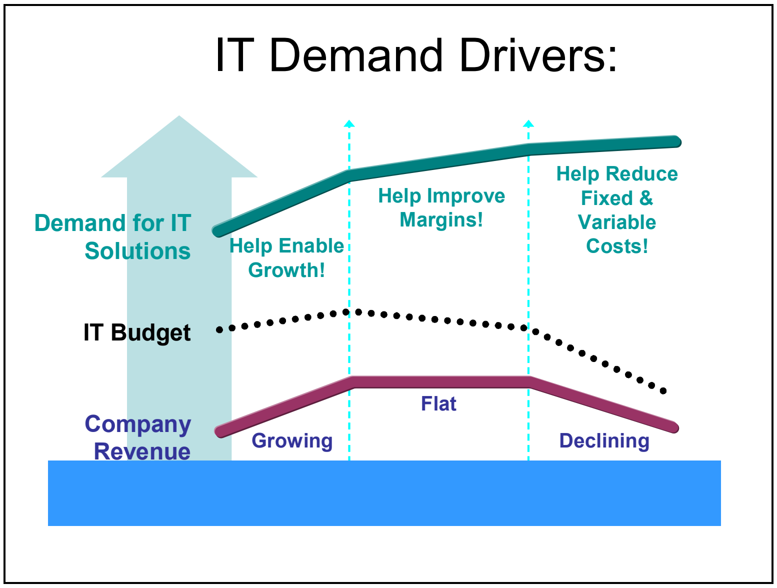

Figure 6-1: Representation of IT Demand Drivers

Server consolidation is one approach to reducing current IT service costs and freeing up IT budget dollars to allow investments in new IT services. Other approaches to reducing IT service costs include improving automation, reducing hardware and software costs through procurement negotiations, and consolidation of facilities.

We are making an assumption that, overall, IT budgets within departments have been, and will trend flat or slightly down. However, new servers are still being added due to new programs or requirements. This puts additional pressure on existing IT staff to manage a greater number of servers with the same head count. Further, the State of California is facing a large rate of attrition, especially in the IT area, and this trend will increase as employees become eligible for retirement. This will also cause an expertise and knowledge drain.

6.2. *Projected Server Growth Rate*

The growth rate for the installed base of U.S. volume servers[7] over the past five years has averaged 12.4% annually. The estimated growth rate for the installed base of U.S. volume servers for the next four years is 9.9% annually[8]. The growth rate for storage is difficult to estimate as it is used in a variety of ways and devices; the demand for storage is estimated to exceed server growth rates.

The historic growth rate of servers for all State of California agencies and departments for the past five years is not available. This study assumes that it was at or below the average U.S. growth rate.

We assume that tight IT budgets will limit the addition of new servers. However, IT budgets will have funds to provide for server refresh. There is evidence in the inventory data that many servers that are more than five years old are still in service. We assume that the installed base of servers is growing by adding new servers for additional applications without refreshing or retiring older servers to a degree. In addition, we assume there is some direct funding for new programs that support adding servers.

Without a special effort at reducing or consolidating the installed base of servers, we project over the next five years that the State will see:

- 5% annual growth rate of installed servers. This is 50% of the estimated 9.9% growth rate of U.S. volume servers and is based on assuming tighter budgets in California Government.
- 15% annual growth rate of online storage.

There is already awareness within the departments of the benefits of server consolidation. The survey results show that several departments have undertaken or are in the process of undertaking server consolidation projects.

With the current effort by departments and a statewide focus on server consolidation, the installed base of servers could be reduced. A manageable three-year goal would be:

- 15% reduction of existing servers.
- 20% conversion of physical servers to virtual servers.

This would flatten the growth rate of servers (physical and virtual) and reduce the number of physical servers. Without flattening the growth rate, additional demand would be placed on existing IT staff to manage the increased number of servers. For an example of this impact, assuming a 9,000 server installed base in 2007, see Figure 6-2: Five-Year Server Growth Rates (5% annual growth rate).

---

[7] IDC - Volume server market (consisting of all systems with an average sales value [ASV] below $25,000)

[8] Derived from IDC Market Analysis: U.S. and Worldwide Server Installed Base 2007-2010 Forecast, IDC # 205504, March 2007.

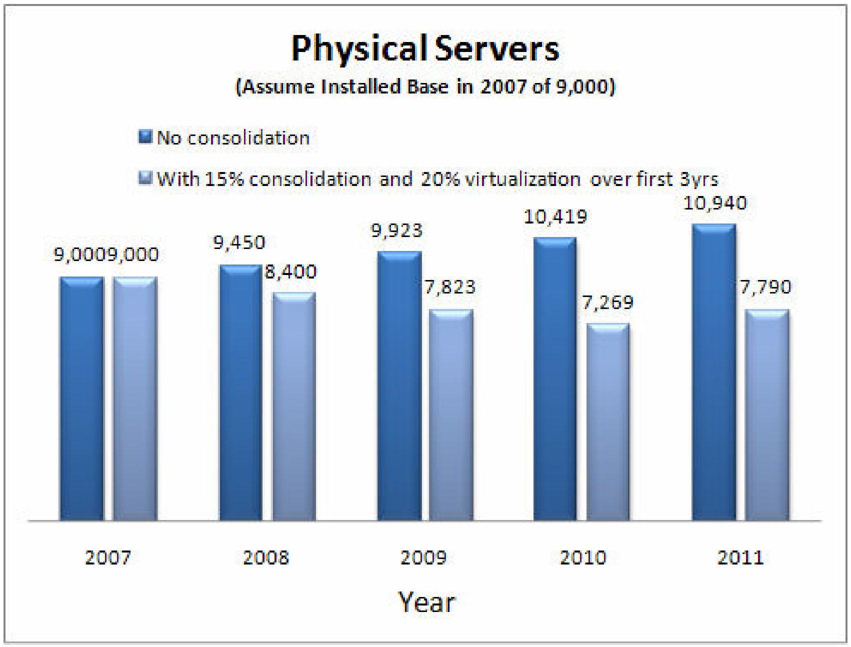

Figure 6-2: Five-Year Server Growth Rates (5% annual growth rate)

7. Current Environment

7.1. State IT Structure

This study focused primarily on the agencies and departments of the Executive Branch of California State Government. State agencies and departments outside the Executive Branch were invited by the State CIO to participate as well.

Departments have independent IT groups that are responsible for meeting the departments' IT needs. This structure forms a federated group of IT departments with common goals and concerns while still leaving operational IT responsibilities and decision making to the departments.

From the IT Strategic Plan[9]:

The California Executive Branch Technology Governance Structure

The IT governance structure will be comprised of the following component layers:

A. A State Chief Information Officer (State CIO) who is a senior advisor to the Governor with full responsibility and authority for statewide technology vision, strategic planning and coordination, technology policies and standards for secure technology solutions, technology architecture, project management and defining a streamlined technology project review and approval process. The State CIO will lead an Office that includes, among other functions, Strategic Planning and Governance, Statewide Policy, Portfolio Management, Enterprise Initiatives, Enterprise Architecture and Workforce Planning.

B. Agency Chief Information Officers (Agency CIOs) who are responsible for overseeing departmental management of assets, projects, data systems, and IT services, through a reporting oversight of departmental CIO's. Each Agency CIO shall develop a 3-year plan to rationalize and standardize within their respective Agency, the IT infrastructure, data, and procedures for all departments within the Agency.

C. A strengthened Departmental CIO function, with Department CIOs directly responsible for all IT activities within the department and accountable to their department director and Agency CIO for purposes of reporting departmental IT performance. All employees in IT classifications and all IT systems, assets, projects, purchases, and contracts will be accountable to the Department CIO, who will, under the direction of the department's governance authority, establish standards and procedures to promote efficient and effective use of IT resources throughout the department. Each Department CIO will develop a 3-year plan to rationalize and standardize the department's infrastructure, data, and procedures,

[9] From "The California State Information Technology Strategic Plan," as updated November 2006.

consistent with the Agency plan developed by the Agency CIO and will report IT performance, accomplishments and issues to the Agency CIO.

Data collection for the study included:

- An inventory of each department's servers,
- Online surveys, completed by the department CIOs and IT staff responsible for servers at each site location, and
- Background research, and targeted interviews. Details concerning the inventory and survey methodology can be found in the Appendix.

The following sections contain our summary observations concerning sites and facilities and the server inventory.

7.2. Sites and Facilities

7.2.1. DTS Data Center Facilities

The Department of Technology Services operates two data centers located within 20 miles of each other in the Sacramento region. The primary data center is referred to as the Gold Camp Campus. Gold Camp is the newer of the two data centers and was commissioned in 1999.

The second data center is the Cannery Campus. This data center is close to downtown Sacramento and was built inside a 1920's-era tomato cannery. There have been numerous upgrades to the Cannery Campus; however, because of its current lease term, old building shell, and security risks, DTS plans to decommission this data center when the lease terminates in 2011. For this reason, this study did not examined the Cannery Campus as a viable site for any centralization aspects of a long-term server consolidation strategy. The successor to this facility will, however, play a crucial role in any long term consolidation/rationalization strategy.

Gold Camp Data Center:

The Gold Camp Data Center is a relatively new facility with many modern capabilities and excellent structural and security characteristics. This data center is a two-story facility with 43,000 ft^2 of raised floor space on the first floor. The second floor is dedicated to general office use.

There is sufficient real estate to expand the raised floor space in this facility by an additional 20,000 ft^2 for a total of 63,000 ft^2.

The reported unused raised floor space in this facility is 6,300 ft^2. However, the study believes[10] that additional space can be gained by minor rearrangement of equipment and removal of the desks and office partitions from the raised floor space.

[10] Based on an on-site walkthrough conducted by one of Intel Solution Services' data-center facilities experts.

The study estimates that this facility in its current state will be able to host between 1,100 and 1,300 standard 2 feet x 4 feet equipment cabinets, if arranged properly.

The power and cooling capabilities of the Gold Camp Data Center are as impressive as the total square footage. This data center has been designed as a Tier III Data Center based on the Uptime Institute tier classification. A Tier III data center is required to be concurrently maintainable. This means that in addition to the required levels of redundancy to protect against failures, all components of the electrical and mechanical systems must be capable of undergoing maintenance or replacement without requiring any scheduled or unscheduled outage to the IT equipment.

Although the Gold Camp facility receives its utility power from a single substation, it has adequate on-site power generation to mitigate any unexpected power failures. Three 1,750 KW (Kilowatt) generators provide a total of 5.25 MW (Megawatts) of total onsite power. However, since one of the generators is a backup generator the total design load cannot exceed 3.5 MW. Two 10,000 gallon diesel storage tanks provide sufficient on-site fuel storage to continuously operate the generators for 5 days.

Downstream of the utility and generators are three 1,100 KVA (~1,000 KW) Uninterruptible Power Supply (UPS) systems which supply power to the IT equipment. There is space allocated to install a fourth system if necessary. The Power Distribution Units (PDUs) and Static Transfer Switches (STS) that are located downstream of the UPS systems are configured such that in the event of a failure or maintenance of a UPS system, the load of the off-line unit can be transferred to the remaining operational units without impacting IT equipment. Because of this redundant configuration, the total design load of the UPS systems is limited to approximately 2,200 KVA (~2,000 KW). The current utilization of the UPS systems is 575 KVA, which is 26% of the total available capacity.

At full load, the UPS systems can provide 15 minutes of battery time to allow for start and synchronization of the generators, which normally require less than 45 seconds.

Three 575-ton chillers satisfy the cooling requirements of the entire Gold Camp facility including office spaces. One of these chillers can cool approximately 2 MW of power consumed by IT equipment, which is equivalent to the maximum load of the data center. There is ample redundancy and capacity in this chiller plant to support loads in excess of 3 MW of power consumed by the IT equipment alone. This is adequate to support installation of the fourth 1,100 KVA UPS system.

In Summary: The project team believes that with minor rearrangement of the raised floor and modification to the airflow dynamics in the raised floor space the Gold Camp data center can play a significant long term role in the State's server consolidation and rationalization strategy. Gold Camp is a relatively new state-of-the-art facility with effective raised floor utilization below 50% and power and cooling utilizations as low as 26%. This leaves ample opportunities to leverage the existing capabilities of this facility. For longer term planning, this facility can be expanded by build-out of another 20,000 ft^2 of raised floor space.

7.2.2. Departmental and Site Facilities

The study requested that each department provide discrete inventories for all servers that were physically resident at each geographically or logically distinct site in its portfolio. Translated, we wanted to see separate inventories for servers at headquarters data-center sites and at each district or local-level remote office.

In some cases, we received seemingly accurate breakdowns of servers by physical site. In many cases, however, it became clear (typically from hints in the server names themselves) that while the inventory was stated as being for one site (typically department-level), it actually contained servers from two or more geographically dispersed sites.

Key Observation: The study cannot conclude anything statistically meaningful from the observation that the majority of departments in the inventory appear to have only one site. Any site-specific observations or recommendations in this document are based on reviews of those few departments for which we have comprehensive site-specific details.

Some departments—namely the Department of Parks and Recreation–have completed a site consolidation of servers. Servers that were once spread across multiple field offices have been centralized and consolidated at the Department's HQ data center. (*It should be noted that this consolidation was only technically possible when sufficient network bandwidth became available to allow now-remote users to obtain satisfactory interactive performance against now-centralized functional servers.)*

Availability of network bandwidth, which was not analyzed in detail in this study, is a key practical consideration in any effort to consolidate. The recommendations described below, to the extent that they involve cross-site consolidation, rest on the assumption that sufficient network bandwidth is obtainable, at an economically viable cost, to support the recommendation.

Key Observation: We do not have reliable site-level information across all departments, but we do have some useful data points. Examination of the details of these sites indicates that many remote sites are equipped with multiple low-end servers, each dedicated to performing a single task. An example would be a directory server, an e-mail server, and a file server. These servers can be consolidated into a single server, with no loss of performance or availability, using either workload consolidation or virtualization.

***Key Observation*:** As we gathered the inventories and conducted selected interviews, we received anecdotal information concerning the current state and future plans for department-level IT facilities. We were told of facilities that have sump-pumps installed on the server floor, because water leaks into the server room when heavy rains occur. We observed extremely crowded, small-scale data centers that developed from available space and accommodate more servers than their design capacity could reliably accommodate. We heard of departments who were planning significant upgrades to existing facilities and/or the creation of new facilities in new buildings to accommodate expansion and/or fix existing problems.

7.3. *Summary of Server Inventory*

7.3.1. Overall Server Inventory Observations

In the inventory collected for this study, there are **6082** servers. Of these, **5753** are physical and **329** are running as virtual machines on other physical servers.

The relatively few AIX (IBM P-Series, for the most part) machines in the inventory are very heavily virtualized (almost 50%). Other platforms are less heavily virtualized on a percentage basis, but Windows 2003 Server represents by far the majority of virtual servers, in terms of absolute count in either category:

Table 6: Servers by OS, Physical vs. Virtual

Phy/Virt	AIX	Free BSD	HP-UX	Linux	MPE/iX	Netware	OS X	Solaris	Tru64	Unixware	Unknown	VMware	W2K0	W2K3	WNT	WXP	Grand Total
Physical	62	2	58	90	78	239	6	232	7	13	628	10	837	3281	138	72	5753
Virtual	56		7	1							25		22	218			329
Grand Total	118	2	65	91	78	239	6	232	7	13	653	10	859	3499	138	72	6082
Percentage	47%		11%	1%							4%		3%	6%			5%

Key Observation: Outside of the AIX environment, server virtualization is low today.

We found virtual servers in use at the following departments:

Table 7: Use of Virtualization by Department

Dept Name	Total
Department of Aging	2
Department of Consumer Affairs	14
Department of Housing and Community Development	16
Department of Insurance	3
Department of Managed Healthcare	9
Department of Motor Vehicles	11
Department of of Corrections and Rehabilitation	82
Department of Technology Services	150
Department of Water Resources	3
Employment Development Department	3
Franchise Tax Board	16
Lottery	20
Grand Total	329

Some departments are using virtualization primarily for development servers, while others are hosting production systems as virtual machines. The following table shows the overall breakdown, by server category, of physical vs. virtual servers.

Table 8: Virtual Server Use by Server Category

Category	Phy/Virt	Total
Development	Physical	642
	Virtual	67
Development Count		709
Production	Physical	4591
	Virtual	227
Production Count		4818
Unknown	Physical	520
	Virtual	35
Unknown Count		555

Key Observation: The State appears to be using virtualization for small-scale workloads. This indicates that there are significant opportunities to use virtualization to reduce server counts by applying the technique to larger-scale workloads also.

7.3.2. Servers by Age

In the inventory, we did not receive information about servers age or install date. We believe that the rated frequency[11] of each server's processor is a suitable proxy for a server's.[12]

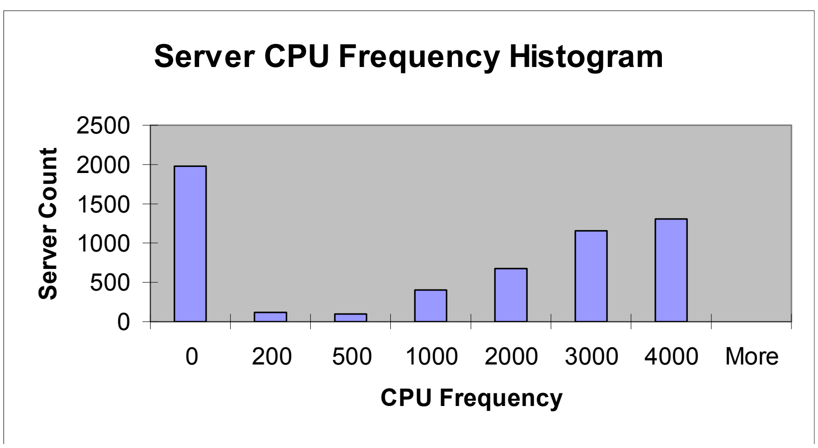

Figure 7-1: Server CPU Frequency Histogram

[11] The number of times the processor's clock cycles each second, expressed in Hz. The performance per clock cycle varies between processor architectures and across generations of the same architecture, but for our use, frequency is a useful first-order measurement that can indicate a servers relative age.

[12] This will not be the case in the future. Future processors will be distinguished by the number of cores in each package instead of per-core processing frequency.

> ***Key Observation:*** If we consider servers with a CPU frequency below 2000MHz as being 'older' and thus potentially good candidates for consolidation via virtualization, then there are at least 618 servers, based on CPU frequency, that are candidates. (Servers with a frequency of 0 are not included, since that value indicates that we have no information concerning frequency for those servers.)

Note that most of the non-X86[13] servers installed at the departmental level are relatively older. We did not analyze the possibility of these servers for virtualization further, but it seems likely that significant consolidation savings could be available from consolidating multiple workloads that are currently running on RISC/UNIX systems onto current-generation RISC/UNIX or X86+Windows/Linux platforms. The latter should be preferred for highly standardized services (e.g. file service, e-mail, directory, etc.).

7.3.3. Servers by Storage Architecture and Capacity

Storage is an important aspect of any consolidation planning effort. Servers that are already using SAN-attached (vs. directly-attached) storage are typically less complex to consolidate.

We are somewhat suspect of the accuracy of the data for this observation, but the inventory indicates that storage-area networking techniques are infrequently applied today. If we focus on just those servers for which we actually have detailed local storage configuration data (~800 of the servers in the inventory), and further focus on just the major functional categories, we find some interesting results (all storage shown in megabytes (MB's)):

Table 9: SAN-Attached vs. Direct-Attached

SAN?	Total
No	5629
Yes	124
Grand Total	5753

[13] The term 'X86' describes all processors, regardless of manufacturer, that are directly compatible with the original Intel-defined architecture first used by the 8086 processors used in the original IBM PC design.

Table 10: Storage Utilization by Primary Function

Main Function	Storage Used vs. Capacity	Total
Application	Average of Storage Used	78,055
	Average of Storage Capacity	175,430
Database	Average of Storage Used	100,498
	Average of Storage Capacity	269,352
Directory	Average of Storage Used	8,355
	Average of Storage Capacity	89,052
Email	Average of Storage Used	159,371
	Average of Storage Capacity	215,692
File/Print	Average of Storage Used	71,691
	Average of Storage Capacity	311,190
Image	Average of Storage Used	129,479
	Average of Storage Capacity	351,300
Unknown	Average of Storage Used	88,955
	Average of Storage Capacity	268,332
Virtual Host	Average of Storage Used	165,274
	Average of Storage Capacity	294,662
Web	Average of Storage Used	21,481
	Average of Storage Capacity	78,605
Storage Used		82,972
Storage Capacity		237,141

Key Observations:

1. In this sample, file servers in particular are over-provisioned in terms of storage. There are many file servers in place, but most of the installed storage in them is underutilized. The main implication of this is that users have a large amount of potentially critical information hidden away on local PC hard drives, where it cannot be easily shared and is more vulnerable to loss and/or theft.
2. Even Database servers, which one would expect to have a fairly large amount of utilized storage, are on average quite small.
3. Only E-mail servers are reasonably well balanced in terms of utilization vs. installed storage capacity. We suspect that, given trends in average message size reported by most commercial organizations, many E-mail servers in the state may find themselves running out of available space fairly soon.

7.3.4. Servers by Primary Function

Most of the detailed consolidation recommendations and the quantification thereof depend heavily on the primary function of the servers being assessed. Specific details describing what we mean by each primary server function are provided in the appendix.

For example, we found many departments with a large number of small-scale file servers co-located at one facility; this is not uncommon. Many commercial organizations have faced this situation. Such file servers have high potential for consolidation, which we describe in detail below, using either virtualization or other techniques (preferred).

The overall breakdown of servers by primary function is as follows:

Table 11: Server Count by Primary Function

Phy/Virt	Application	Backup	Communications	Database	Directory	Email	File/Print	Gateway	Image	License	Management	Remote Access	Security	Unknown	Virtual Host	Virus	Web	Grand Total
Physical	1481	122	12	607	404	365	882	106	14	4	215	84	169	914	28	20	326	5753
Virtual	87	1		46	26	11	28	1	1		8	6	15	61		2	36	329
Grand Total	1568	123	12	653	430	376	910	107	15	4	223	90	184	975	28	22	362	6082
	6%	1%	0%	7%	6%	3%	3%	1%	7%	0%	4%	7%	8%	6%	0%	9%	10%	5%

The greatest opportunities for reducing server counts (and therefore operational and maintenance expenditures over time), will come from applying best known consolidation methods to the **Application, File/Print, Database, Directory, E-mail,** and **Web** categories.

Virtualized server instances are shown in this chart to indicate the intensity with which virtualization, as a form of consolidation, has already been applied within functional categories. It is not surprising that the Directory and Web categories show relatively high rates of virtualization – these are often the least utilized servers in any organization's infrastructure, and are among the best candidates for virtualization.

What is somewhat surprising is the extent to which Database servers have been virtualized. We suspect that this is in part a function of the amount of virtualization already in place on AIX systems within DTS, and in part a result of the fact that many database servers are deployed on very low-end, probably underutilized systems, making them good candidates for virtualization.

In the remainder of this chapter, we will limit the discussion to higher-level descriptive attributes for the servers in each of these categories. The figures cited in the following sections will be for physical servers only. Virtual machine instances have been filtered out.

7.3.5. Database Servers

Database servers can be challenging to consolidate if they are very large. On a number of dimensions, however, it appears that there are relatively few 'large' database servers in use in the State. Using CPU count as a proxy for size, we see that most database servers are concentrated in the 4-processor and below category[14]:

[14] Note that processor count may be over-reported in some cases for Intel-based servers installed during the last 5 years. Many Intel processors during this period were equipped with a 'Hyperthreading' feature, which appears to the operating system as if it is an independent processor.

Table 12: Database Server CPU Speed by CPU Count

CPU Count	Attributes	Values
1	Average of CPU Speed Server Count	1796 72
2	Average of CPU Speed Server Count	2248 331
3	Average of CPU Speed Server Count	180 2
4	Average of CPU Speed Server Count	2570 135
5	Average of CPU Speed Server Count	440 1
8	Average of CPU Speed Server Count	2212 28
12	Average of CPU Speed Server Count	1200 1

This table shows the number of database servers grouped by the number of CPU's available for each server. Additionally, it shows the average CPU speed (frequency) of all servers in each CPU-count bucket. Our observations indicate that RISC systems outside of DTS tend to be quite old. We did not analyze the possibility in depth, but there could be substantial savings available from retiring old RISC systems and migrating their workloads to more modern platforms. Unless there is some specific consideration[15] that requires the use of a proprietary RISC/UNIX platform, migrating to a modern X86 platform running either Windows or Linux should be the preferred approach.

Key Observation: Many database servers in the State are good candidates for consolidation, using either virtualization or the more elegant approach of database instance consolidation on fewer, larger-scale dedicated database servers.[16]

15 Like a particular version of a DBMS, particularly one that is no longer available or supported, which is required by an application.

16 This approach can be problematic when application requirements dictate that multiple different DBMS versions remain available. Usually a combination of virtualization and instance consolidation is necessary to achieve maximum consolidation ratios.

7.3.6. File Servers

The file server population is dominated by X86 architecture, unsurprisingly:

Table 13: File Servers by Platform Architecture

CPU Mfr	Data	Total
Sun	Count of Servers	17
	Average of CPU Speed	200
AMD	Count of Servers	1
	Average of CPU Speed	2000
IBM	Count of Servers	1
	Average of CPU Speed	400
Intel	Count of Servers	802
	Average of CPU Speed	1792
Unknown	Count of Servers	61
	Average of CPU Speed	0
Total Count of Servers		882
Total Average of CPU Speed		1636

Table 13 clearly shows that there are instances of old proprietary RISC/UNIX systems in use that bear replacement—see the Average CPU Speed for Sun and IBM platforms. For basic shared-file services, we suggest using commodity X86 servers.

Key Observation: One of the key considerations for file server consolidation is geographic distribution. As noted previously, the quality of the inventory data in terms of site-level server distribution is poor. So our analysis will be based on a theoretical, instead of data-driven, approach.

7.3.7. E-mail Servers

Like Database and File/Print, most e-mail servers in use across the State are based on X86 architecture and are fairly small:

Table 14: E-mail Servers by CPU Count and Speed

CPU Count	Data	Total
0	Count of Servers	25
	Average of CPU Speed	0
1	Count of Servers	33
	Average of CPU Speed	1997
2	Count of Servers	234
	Average of CPU Speed	1715
4	Count of Servers	64
	Average of CPU Speed	2513
8	Count of Servers	6
	Average of CPU Speed	2519
16	Count of Servers	3
	Average of CPU Speed	3000
Total Count of Servers		365
Total Average of CPU Speed		1786

The existence of a small number, exactly three, of very large-scale e-mail systems indicates that a degree of e-mail consolidation is already underway. All of these large-scale e-mail servers are located at the DTS data centers[17].

Key observation: Over 250 e-mail servers in the inventory are good candidates for consolidation based on their observed scale, either at the local site or centrally.

7.3.8. Web Servers

As with other server categories, the Web server population is dominated by relatively low-end X86 architecture-based systems. The average CPU speed of the 183 2-way systems dedicated to Web serving further indicates that these systems are rather dated:

Table 15: Web Servers by CPU Count and Speed

CPU Count	Data	Total
0	Count of Servers	44
	Average of CPU Speed	0
1	Count of Servers	44
	Average of CPU Speed	2430
2	Count of Servers	183
	Average of CPU Speed	1412
4	Count of Servers	48
	Average of CPU Speed	2554
8	Count of Servers	7
	Average of CPU Speed	2793
Total Count of Servers		326
Total Average of CPU Speed		1556

The State, like many private organizations, suffers from the after-effects of the ‘Webification of everything’. With the implementation e-government and e-business, there was rapid deployment of web sites and their attendant web servers. An examination of the details of the Web servers themselves reveals that there many different approaches to Web site development.[18]

Key Observation: Web servers are frequently good candidates for centralized consolidation. Our anecdotal observation from examining many of the websites maintained by departments that submitted inventories is that many sites use static Web pages. Therefore, there are no practical network-bandwidth based barriers to centralized consolidation of many Web servers, and there are potentially many benefits thereof.

[17] Note that these ’16 way’ servers are very likely to actually be 8-processor servers with Hyperthreading enabled.
[18] This comment is based on an examination of the details of the run-time and development-time software stacks installed on various Web servers across multiple departments.

7.3.9. Application Servers

Application servers are more difficult to analyze. They are also more difficult to consolidate through any method other than virtualization. Additionally, application server cross-site consolidation is typically <u>even more heavily dependent on network bandwidth</u> considerations than infrastructure server consolidation. However, there still appear to be significant opportunities for consolidation of application servers:

Table 16: Application Servers by CPU Count and Speed

CPU Count	Data	Total
0	Server count	239
	Average of CPU Speed	0
1	Server count	304
	Average of CPU Speed	1391
2	Server count	773
	Average of CPU Speed	2123
3	Server count	5
	Average of CPU Speed	1388
4	Server count	143
	Average of CPU Speed	2467
5	Server count	1
	Average of CPU Speed	0
8	Server count	11
	Average of CPU Speed	2512
12	Server count	1
	Average of CPU Speed	180
16	Server count	4
	Average of CPU Speed	2993
Total Server count		1481
Total Average of CPU Speed		1663

As with other categories, most application servers are based on relatively low-end server configurations. If industry-typical utilization rates of 15% or less[19] are relevant for the application servers in use around the State[20], then the application server category is a large potential source for consolidation. This is true even if application server consolidation is limited to consolidation within the sites where these servers are located today, <u>with no consolidation to a central facility assumed</u>.

Key Observation: While application servers have high consolidation potential, realizing that potential can be quite challenging, especially when cross-site consolidation (e.g. centralization) is pursued. It was beyond the scope of this study to gather sufficiently detailed information about application servers to provide meaningful recommendations for application servers other than for on-site consolidation via simple virtualization techniques.

[19] Source: IDC, September 2006, server virtualization study.
[20] We did not measure actual resource consumption for any servers as part of this study.

8. Key Recommendations

We look at five key consolidation recommendations in this section:

- Data Centers
- In-Department Consolidation
- File Sharing and Content Management
- E-mail
- Virtualization

The cost and value of each of these recommendations is dependent on architecture, design selected, and implementation details. We provide an illustrative analysis of cost and value based on the server inventory and survey results collected as part of this study. Actual savings are dependent on implementation.

We have not factored into our analysis the cost of infrastructure upgrades, such as the WAN network. These are real costs, but items like network upgrades benefit considerably more than just server consolidation, such as collaboration, VoIP, and rich media. WAN networks have advanced and will continue to advance, providing even more opportunity for advanced applications such as mobility, video, and multi-media training, just to name a few.

We identify risks as part of the analysis. While there are risks associated with these consolidation recommendations, we feel that they can be managed and mitigated. Management sponsorship and support along with good project planning is required for success.

These recommendations are derived from a high-level review. As a next step, it is important to complete a further detailed analysis of each recommendation to fully understand their impact.

Note: There is overlap between those servers that would be eliminated in the In-Department Consolidation scenario and the File Sharing and E-mail scenarios. For this reason, it cannot be assumed that the cost and savings are additive across all of the key recommendations.

8.1. Data Centers

The question of server room and data center (referred to here as computer room) consolidation and centralization came up frequently and is a controversial subject for several departments. A full discussion of the subject extends beyond this study and the implications of consolidation. Several organizational roles and responsibilities issues need to be considered and addressed as part of the broader subject.

For this study, we consider the importance of computer room utilization in the efficient use of resources. While we understand there are further organizational issues to be considered and technical issues to be worked out, there is substantial benefit in some consolidation of computer room facilities and the setting of future facilities-use guidelines.

8.1.1. Recommendation

The Gold Camp data center is significantly underutilized. A second facility is planned for co-processing, business continuity, and to replace the Cannery data center. No significant addition or expansion should be made to existing departmental computer room facilities. Rather, the Gold Camp and Cannery successor data center facilities should be utilized. Based on State developed minimum requirements, current computer room facilities should be evaluated to determine if they meet the minimum requirements for security and continuity of operations commensurate with the applications and server functions they support. If not, the processing should be moved to the DTS data centers.

8.1.2. Detail

Note: This recommendation addresses relocating servers and providing for support of the hardware by DTS. It does not address a change in the support of the operating system or applications (software). In addition, network access by end users is identified as a potential risk but it is assumed that this risk can be addressed and mitigated.

There are a number of computer room facilities within the State's departments. These range from very small server room closets that host just a few servers to large, fully equipped data centers. In the survey completed as part of this study, there were 48 site surveys completed and 55 computer rooms reported (computer rooms are housed at departmental sites). The survey responses represent only a portion of the total sites and computer rooms. The survey data also includes remote sites that have small server rooms to support local processing and remote offices.

The study did not evaluate the utilization, efficiency, or security of these computer rooms. The State should develop a minimum data center, computer room, and server room standard for the various types of servers in use at the State. A review of computer rooms should be done to ensure they meet minimum standards. Those that provide efficiency, meet minimum standards, and do not require significant upgrade are not part of the recommendation to move to the DTS data centers.

The State has invested significantly in the Gold Camp data center. It currently has capacity (floor space and electric) for at least 4,000 – 5,000 additional midrange servers, supporting networks, and data storage. Part of the data center floor space is used by offices that can be moved. The facilities master plan also allows for expansion. This does not consider moving servers and equipment from the Cannery to Gold Camp or capacity at a future second data center.

An offering of DTS allows for "Customer Owned Equipment Managed Services (COEMS)." The implementation of this offering should be reviewed to make sure that it does not prevent the full utilization of the data center facilities. We recommend that the DTS staff manage all hardware at the data center. This would reduce the need for department staff having to travel to the data center and face delays. DTS would provide full hardware support and remote hands[21] service when needed.

For many departments, utilizing DTS data centers could simplify business continuity planning as a second facility for business continuity is in place and existing DTS processes and procedures can be used.

8.1.3. Cost and Value

The cost savings / avoidance to the State are unknown but presumed to be substantial over time. The potential areas for cost savings / avoidance are:

- Using the existing DTS Gold Camp data center facility can improve utilization and efficiency.
    - Increasing utilization at the data center can be done with minimal additional facilities costs.
    - The underutilization at the data center causes the allocation of the facilities, resources, and staff to be spread over a reduced base of servers. If data center utilization was increased, the per server cost for facilities and services would be reduced substantially.
- Consolidating hardware support personnel builds better expertise across a reduced staff.
- Retrofitting existing departmental computer rooms to meet minimum standards would be costly.
- Building new facilities at one or more departments would be costly.

A State minimum standard for data centers, computer rooms, and server rooms needs to be created. We assume for security, efficiency, and continuity of operations that Tier II or Tier III facilities should be required for the majority of State production servers and data.

As an example, the cost to construct a modern Tier III data center can be estimated using the following:

- Assume $220 per sq. ft. of raised floor space
- Add electrical and mechanical cost of $20,000 per Kilowatt of UPS power
- This does not include annual maintenance and operating costs

[21] "Remote hands" is an IT industry term to mean someone who can do something for someone else who is remote, e.g. reboot a server, look at the actual server monitor, insert a CD, enter commands during server startup that cannot be done remotely.

Therefore, the cost of constructing a data center facility similar to Gold Camp, excluding land and office space is over $49M as shown in the following table.

Table 17: Tier III Data Center Construction Costs

Item	Unit	Unit Cost	Total
Raised Floor Space (sq. ft.)	43,000	$220	$9,460,000
Electrical/Mechanical (KW)	2,000	$20,000	$40,000,000
TOTAL			**$49,460,000**

Also for comparison, a recent review at Intel Corporation in support of server consolidation within their IT department estimated the cost of building a new data center module to support ~4,400 servers (of currently average deployed size) would cost ~$30M.

Other benefits:

- Improved business continuity / Continuity of Government
- Reduced power requirements through improved power management
- Better inventory control of hardware assets
- Ability to negotiate better vendor hardware support rates as equipment will be in one place
- Improved mean-time to fix as spare hardware and parts can be readily available
- Maintain hardware refresh rates, thereby reducing failure rates and benefit from improved performance and energy savings
- Unified Storage through DTS Enterprise Storage implementation.
- Ability to deploy and manage a farm of physical hosts for virtual machines from multiple departments

8.1.4. Risks

There are risks that need to be addressed related to consolidating servers at the DTS data centers. In addition, several issues and concerns are presented in Chapter 10 - Issues and Concerns and should be reviewed:

Risk	Mitigation
Cost of implementation – moving servers and/or services requires planning and execution costs.	• A detailed project migration plan must be in place. • Infrastructure upgrades need to be planned. • It is assumed that most of these moves would be in lieu of a server room retrofit or new construction, which would also require project planning, and execution costs. • The costs would be determined as part of the planning process.
Network – to move servers and/or services to the DTS data centers, there	• If sufficient WAN bandwidth does

must be sufficient network bandwidth and redundancy to support a remote user base. Please see Chapter 10.3 - Issues and Concerns, Network.	not exist, it can be acquired. • Network cost to the department is a risk and a concern for some departments. These costs need to be addressed but should be offset by cost avoidance of building or retrofitting facilities. • Metropolitan Area Network (MAN) architecture should also be considered. DTS already provides some MAN. • As part of the project, a detailed migration plan should be developed.
DTS rates – the requirement for how DTS receives funds and recoups costs is set by policy. This policy needs to be reviewed and revised to not penalize first adopters for new services and the underutilization of facilities and services.	• DTS rates need to be reviewed. • Department's costs to run a server or service should include all applicable costs, including those for refresh, facilities, security, DR, and staff costs. • DTS should be able to provide service at comparable and reasonable rates to what departments have today.
Incorrect Determination of Data Center Capacity – if the assessment of the available capacity in this study is not correct then the ability to centralize servers will be limited.	• A further assessment of the capacity of each DTS data center needs to be completed. • Plans for new DTS data center facilities need to be completed.

8.2. *In-Department Consolidation*

While data center consolidation provides the largest cost avoidance and savings, it does not reduce the number of physical servers. The best opportunity for server consolidation remains with the departments. Server consolidation has been an industry initiative and best practice for several years and several State departments have completed some server consolidation or are consolidating servers now.

8.2.1. Recommendation

The State CIO should set a goal to eliminate 15% of existing servers through combining workloads and services over the next three years. In addition, a plan should be developed with the cooperation of the departments to meet this goal. A ***simple*** quarterly tracking spreadsheet/system should be set up to record information by department. The system could track the total number of servers, the number of physical servers, the number of virtual servers, the number consolidated during the quarter, and the number virtualized during the quarter. The spreadsheet and progress should be reported quarterly to the State CIO.

8.2.2. Detail

In-department consolidation does not preclude housing servers at the State data centers. While the management of server functions and applications can remain with the department, the physical location of servers for this recommendation is not important.

This recommendation is for the elimination of servers altogether and not just converting physical servers to virtual servers. A further recommendation to migrate physical servers to virtual servers is presented below in Section 8.5: Virtualization.

A goal to eliminate 15% of the exiting servers over a three-year period can be challenging but should be manageable. This goal could include e-mail servers and file servers that are also a separate recommendation. Other server consolidation opportunities are presented in Chapter 9: Additional Opportunities.

Combining Workloads:

Departments have the greatest insight into their business needs, business applications, and their suitability for consolidation.

There are more potential cost savings if an entire application or instance of a server can be eliminated. The largest factor in TCO is support. By eliminating entire applications, the associated hardware, operating system, and application support costs can be reduced or eliminated. Departments should look for applications that support a limited user base or that have overlapping functionality with other applications.

After looking for applications that can be eliminated, the next opportunity is applications running on separate servers that can be combined, thereby eliminating hardware and operating systems and their support. This is referred to as ‘shared landing’. Some opportunities may exist that are relatively easy to take advantage of, but in general, this is a difficult consolidation opportunity. Application vendors should be encouraged to develop and support their applications in this shared landing approach.

Good candidates for combining workloads include:

- Application servers
- Database systems
- Web servers
- Terminal Services

Departments can also use the suggestions for servers to consolidate from Chapter 9: Additional Opportunities.

Combining workloads in a single system has a greater potential for savings than server virtualization.

Quarterly Tracking:

The question of statewide server consolidation is raised in various places, including at the State CIO's office, the Legislature, State agencies, and State departments. Since there have already been efforts within the departments to consolidate servers, these efforts and results should be tracked and highlighted.

For this initiative to be successful, departments need participate in developing a plan and reporting progress. A ***simple*** quarterly tracking spreadsheet/system can be implemented to record information by department. The system could track the number of total servers in each department, the number of physical servers, the number of virtual servers, the number consolidated during the quarter, and the number virtualized during the quarter. The spreadsheet and progress should be reported quarterly to the State CIO. This would retain focus on the goal and highlight progress being made.

Existing Consolidation Efforts:

Through the server inventory, the survey, and interviews, we learned of several departments that have actively engaged in server consolidation efforts or have consolidation efforts in progress. These efforts should be recognized. In addition, through the IT Council, experiences and best practices can be shared.

The following departments reported at least some server consolidation effort in the survey:

Table 18: Servers Already Consolidated By Department

Department	Servers Reported Consolidated
CA Environmental Protection Agency (Responding for Ca Integrated Waste Management Board also)	5
California Energy Commission	5
California State Library	1
Department of Corrections and Rehabilitation	49
Department of Finance	13
Department of Fish & Game	Beginning
Department of Forestry & Fire Protection	10 (plan for 50%)
Department of Housing & Community Development	NR*
Department of Industrial Relations	20
Department of Justice	NR*
Department of Managed Health Care	8
Department of Parks and Recreation	10
Department of Pesticide Regulation	15
Department of Technology Services	90
Fair Employment & Housing Commission	2
Franchise Tax Board	50
Office of Emergency Services	7

* NR – Not reported

8.2.3. Cost and Value

This cost and value analysis is a "What if?" scenario, assuming a 15% reduction (elimination) of existing servers. A scenario for migrating physical servers to virtual servers is covered in the Virtualization Key Recommendation discussed in Section 8.5: Virtualization.

The inventory collected 6,000+ servers. The total number of servers across the State is unknown, but we assume there are over 9,000. The project team used Microsoft Windows and Intel-based dual-socket servers as the proxy[22] for servers but there is also opportunity for consolidation of larger Windows and UNIX servers. The scenario does not phase in the elimination of servers over time but is intended to show the overall cost reduction; the actual cost reduction would be gradually realized across several years.

Assumptions:

- There is an installed base is 9,000 servers.
- 15% of this installed base—1,350 servers—will be eliminated by combining applications onto fewer systems.

[22] Proxy – used to substitute for the variety of servers found in the inventory.

- Windows Server 2003 and Intel 2-processor[23] systems will be used as a proxy for the servers being eliminated.
- This scenario does not account for any servers being virtualized.
- The elimination of servers is all at once.
- Project implementation cost of $2.5M—includes consultant fees, solution design, and implementation. Specific project implementation costs are highly variable, are subject to numerous factors, and design decisions. We supply a judgment cost here to provide a more accurate analysis.
- The FTE[24] to server ratio used for generic servers is 1:15.

Environment

The server configuration priced in the TCO analysis for the In-Department Consolidation scenario is: 2-way Xeon Core 2 Duo processors, 2 GB of memory, SAN attached.

NOTE: The same server configuration is used as a proxy for servers in the current and the recommended scenarios. The server configuration approximates the most common specifications of servers reported in the server inventory.

Table 19 shows the major components of the TCO model priced in this study.

	Current In-Department Model	*Recommended In-Department Model*
Number of physical servers:	9000	7650
Number of virtual servers:	0	0
Number of software licenses:	9000	7650
Number of FTEs:	600	510
Data center floor space (square feet):	38621.7	32828.4
Project Implementation:	0	$2.5M

Table 19: Major TCO components

[23] When we use the word ‘processor’ in this context, we’re really talking about ‘sockets’, since basically all modern server processors include multiple processing cores per physical socket on the motherboard.
[24] FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance. FTE’s by workload are based on CIOView standard metrics.

TCO Analysis Results

Table 20 shows the estimated cost per year of the current server environment as modeled in this TCO.

Current In-Department	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	6,184,000	6,184,000	6,184,000	8,374,000	8,374,000	35,299,000	9.37%
Software	4,500,000	4,500,000	4,500,000	4,500,000	4,500,000	22,500,000	5.97%
Facilities	10,617,000	10,617,000	10,617,000	10,617,000	10,617,000	53,086,000	14.09%
Personnel	53,180,000	53,180,000	53,180,000	53,180,000	53,180,000	265,902,000	70.57%
Implementation	-	-	-	-	-	-	0.00%
Total Cost of Ownership	74,481,000	74,481,000	74,481,000	76,671,000	76,671,000	376,786,000	100.00%

Table 20: Estimated total cost of current server environment

Table 21 shows the estimated cost per year of the in-department consolidation solution. The implementation costs include disposing of 1350 old servers, and a $2.5M charge for project implementation spread over three years.

In-Department Opportunity	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	5,256,000	5,256,000	5,256,000	7,118,000	7,118,000	30,004,000	9.29%
Software	3,825,000	3,825,000	3,825,000	3,825,000	3,825,000	19,125,000	5.92%
Facilities	9,025,000	9,025,000	9,025,000	9,025,000	9,025,000	45,123,000	13.98%
Personnel	45,203,000	45,203,000	45,203,000	45,203,000	45,203,000	226,016,000	70.01%
Implementation	892,000	833,000	833,000	-	-	2,558,000	0.79%
Total Cost of Ownership	64,201,000	64,142,000	64,142,000	65,171,000	65,171,000	322,827,000	100.00%

Table 21: Estimated total cost of in-department consolidation recommendation

Table 22 shows the potential raw savings numbers—for each category across five years—as the difference between the cost of the in-department consolidation recommendation and the current environment. The total potential savings is about $54M over five years.

Value of Opportunity (Cost shows as negative)	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware	928,000	928,000	928,000	1,256,000	1,256,000	5,295,000
Software	675,000	675,000	675,000	675,000	675,000	3,375,000
Facilities	1,592,000	1,592,000	1,592,000	1,592,000	1,592,000	7,963,000
Personnel	7,977,000	7,977,000	7,977,000	7,977,000	7,977,000	39,886,000
Implementation	(892,000)	(833,000)	(833,000)	-	-	(2,558,000)
Total Savings	10,280,000	10,339,000	10,339,000	11,500,000	11,500,000	53,959,000

Table 22: Value of in-department consolidation recommendation

Table 23 depicts the total cost and savings per category.

Costs	Current In-Department	In-Department Opportunity	Savings	Value	% Savings
Hardware	35,299,000	30,004,000	Hardware Savings	5,295,000	15.00%
Software	22,500,000	19,125,000	Software Savings	3,375,000	15.00%
Facilities	53,086,000	45,123,000	Facilities Savings	7,963,000	15.00%
Personnel	265,902,000	226,016,000	Personnel Savings	39,886,000	15.00%
Implementation	-	2,558,000	Implementation Savings	(2,558,000)	N/A
Total Cost	$ 376,786,000	$ 322,827,000	Total Savings	$ 53,959,000	14.32%

Table 23: Total cost and savings per category

The savings in every category are 15% because we are reducing the server count by 15%, without replacing old servers with new servers. The total savings is reduced to 14.32%, $54M, due to the cost of disposing of the old servers and a project implementation cost of $2.5M.

Figure 8-1 compares the total costs per year of each environment. The savings are about 15% every year for five years.

Summary	Year 1	Year 2	Year 3	Year 4	Year 5	Total	Savings
Current In-Department	74,481,000	74,481,000	74,481,000	76,671,000	76,671,000	376,785,000	53,958,000
In-Department Opportunity	64,201,000	64,142,000	64,142,000	65,171,000	65,171,000	322,827,000	

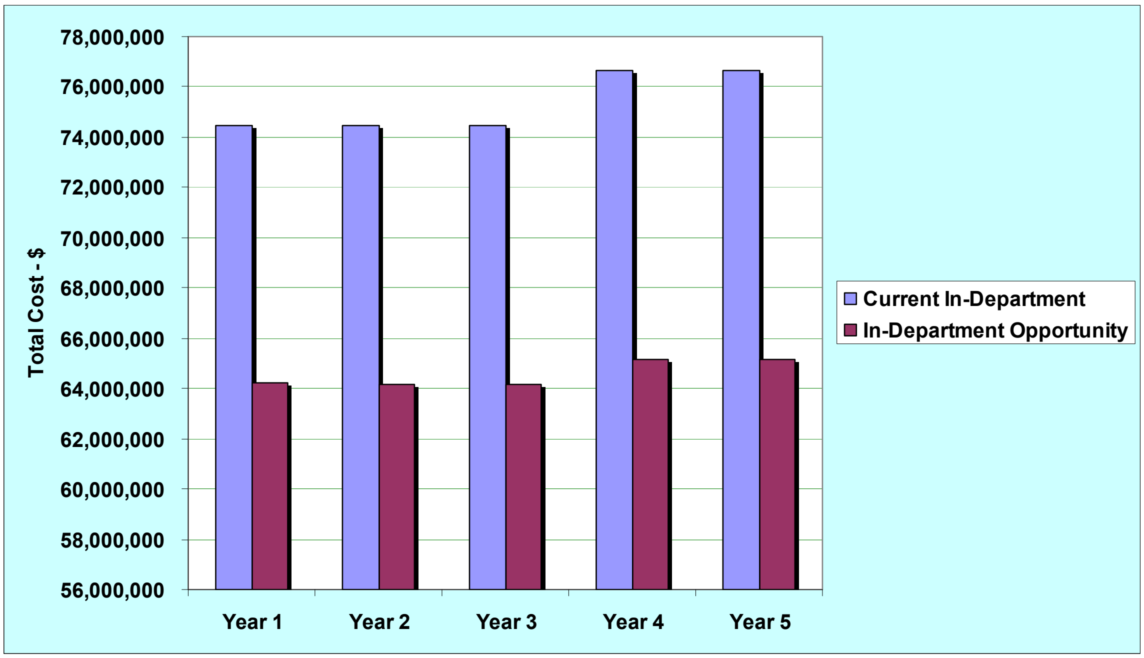

Figure 8-1: In-department consolidation recommendation total cost comparison

Key Financial Result:
The estimated total savings from consolidating the workloads of 9000 servers to 7650 servers is $54M, or 14.3% over five years.

Energy Conservation:

By disposing of 1,350 physical servers, the State can conserve 7,899 megawatt-hours per year. This savings translates to a 15% reduction in energy consumption relative to the current In-Department environment presented in this "What-if?" scenario.

Other benefits:

- Reduced installed server base can reduce workload of existing IT staff.
- Reduction of facilities requirements may extend life of current facilities.

- Funds and resources can be applied to other projects.

8.2.4. Risks

There are risks that need to be addressed related to in-department consolidation. Also, review issues and concerns that are presented in Chapter 10 - Issues and Concerns:

Risk	Mitigation
Complexity – merging or retiring applications can be complex.	• A detailed project plan must be in place. • The costs would be determined as part of the planning process.
Competing Priorities	• The priority of consolidation servers and applications has to be compared with other projects. • The long-term benefit need to be appreciated as part of project evaluation.

8.3. *File Sharing and Content Management*

The detailed chart in the appendix labeled **Error! Reference source not found.** shows the count of virtual and physical file servers for each site in the inventory.

A close examination of this and other data supports the following observations:

- The ratio of file servers to total servers is highly variable from site to site, even within departments.
- On average, file servers in use in the state are small-scale in terms of compute, memory, I/O and storage capacity, at least when compared to other server types.
- Even given the limited storage capacities of most file servers, installed storage is typically less than 50% utilized.
- The only file servers in the State that were identified as being connected to a shared SAN storage utility were those at the Department of Motor Vehicles. All others used direct-attached storage.

We did not gather detailed information concerning the number of users connected to file servers, the geographic locations of these users (e.g. – their 'network distance' from their primary file servers), or network saturation levels.

Key Observation: Small amounts of file-server storage with low utilization are typically a strong indicator that file servers are not being used as the 'repository of record' for departmental documents. It is likely that such documents are being stored primarily as attachments to e-mails and/or as files on local PC hard drives. If so, there are at least two issues: 1) such documents will be hard to find, and 2) recovery in the event of disaster will be extremely difficult.

8.3.1. Recommendations

Near-Term: Where practical, sites with more than two co-located file servers should review utilization and consolidate these servers to two clustered file servers.

Strategic: Evaluate the potential for applying Wide-Area File Systems (WAFS) technology for remote sites. In addition, conduct an analysis of the costs and benefits of implementing a statewide Enterprise Content Management (ECM) service that can provide a more robust and capable document management capability.

8.3.2. Details

Statewide File Server Architecture:

The file server architecture in use around the State today is quite conventional. Few file servers are set up in high-availability configurations. If a file server dies, end users will be unable to access their files until the server is brought back online.

Simple File Server Consolidation

The simplest consolidation approach for file servers, architecturally, is shown in this diagram:

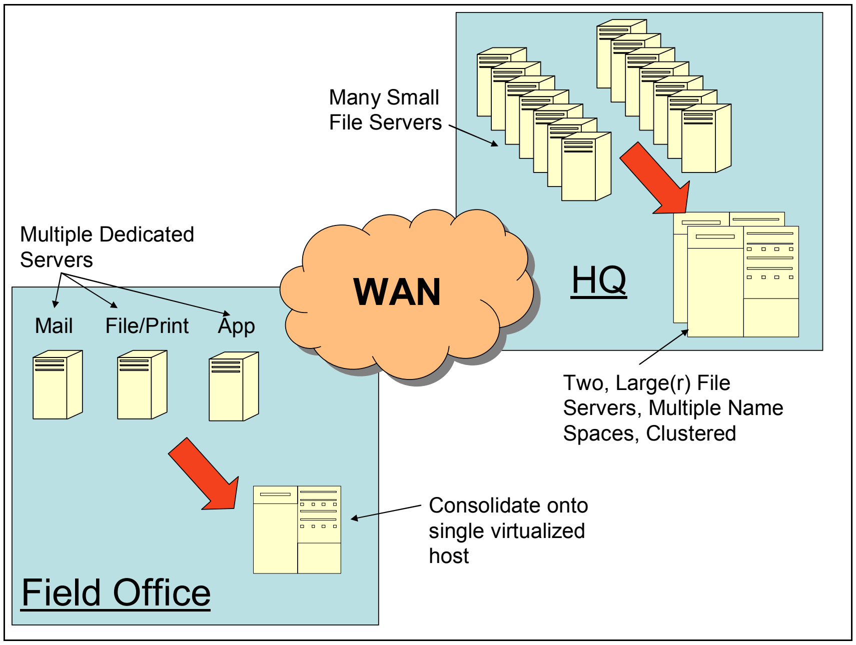

Figure 8-2: Basic File Server Consolidation Architecture

This architecture will result in a net reduction in the number of file/print servers in use across the state. However, it does not make any fundamental improvements to things like cross-site access to frequently used files, high availability for remote sites, backup/recovery operations, etc.

Wide-Area File System-Based Consolidation

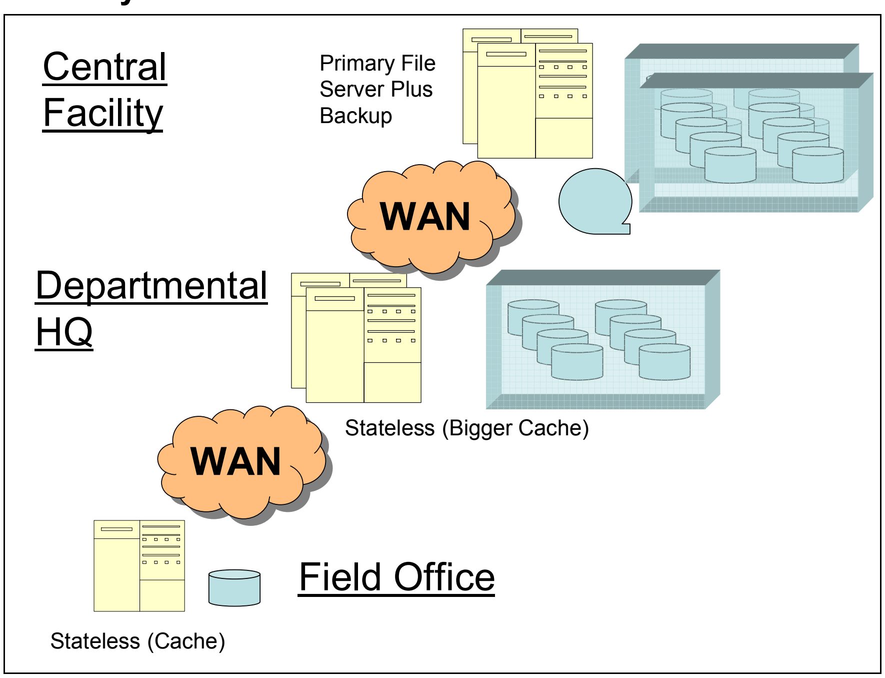

Figure 8-3: Wide-Area File System Architecture

Wide-area File System (WAFS) technology eliminates most of the management tasks associated with dispersed conventional file servers, without affecting local control over content access rules. As far as the end user is concerned, at any level of the distribution hierarchy, nothing changes with WAFS. Users still see the same file servers, redirected drive letters, and name spaces that they have always seen.

The difference is that the local file server, which serves their shared data, does not really contain all of the data within the file system – it only contains a (tunable) portion of the overall file system – typically only the largest, most frequently used files. Software running on each server node in this architecture detects when the user has requested a file that is not already stored locally, and retrieves the requested file from the next-level server in the hierarchy. System managers create policies that govern the tradeoff between file-access latency, local disk storage, and network bandwidth utilization that are implicit in this architecture.

Benefits:
- **Eliminates the need for local backup operations** – only the 'copy of record' of all files stored at the highest level of the hierarchy requires backup.
- Protocols are **optimized for lower-bandwidth WAN links**, including sporadically connected WAN links. The result is much lower-latency access to network-remote files than would be the case with classic NAS protocols.
- **Inherent high-availability** – if the local cache server instance fails, then the next server level up in the hierarchy can provide file service, transparently. When the local server comes back up, local users see a performance improvement, but they have never lost access to their files during the outage.
- **Facilitates centralized file searches**. If a search that spans the name spaces of multiple departments needs to be conducted, it can be run at the central repository, with a guarantee that all information there is very nearly current (only in-flight updates would be omitted from the search). There is no need to go out over the WAN to retrieve and search files that are only stored locally – they can be accessed at the central site at full speed.
- **Reduces overall storage requirements**. WAFS solutions generally have the effect of reducing the overall amount of storage required, since they eliminate or at least control the amount of file duplication that is occurring.

Limitations:
- **Limited metadata**. The WAFS approach is still a file server. As such, it offers no advantages for finding files other than the above-described centralized search benefit.
- **Increases demand for network bandwidth and reliability**. If a file is not found locally, and must be retrieved from a higher level in the hierarchy, response time is delayed if the network is congested or unavailable.
- **Increases complexity**. While local administration of access rights within name spaces does not necessarily change, WAFS demands a new administrative function – set up and administration of the links between the leaf nodes and the higher levels in the hierarchy, as well as for the central site.

Key Observation: Implementing a WAFS approach could be a useful (and re-usable) near-term implementation step in the process of implementing a comprehensive statewide enterprise content management (ECM) service. The servers used in the near term to deploy WAFS could be re-purposed to support the document caching functionality that will be required to implement a high-performance ECM service.

Enterprise Content Management Replacement

Enterprise Content Management (ECM) architectures do not have user-visible file servers. Instead of navigating a directory/subdirectory/file-oriented access method, ECM users access documents using a variety of query-oriented interfaces.

These interfaces, which can be Web-based or native, allow the user to search the document base on virtually any attribute – not just the file name.

When the user finds the document they wish to read or modify, depending on access rights in the document management repository, they tell the interface to open the document, and the document appears.

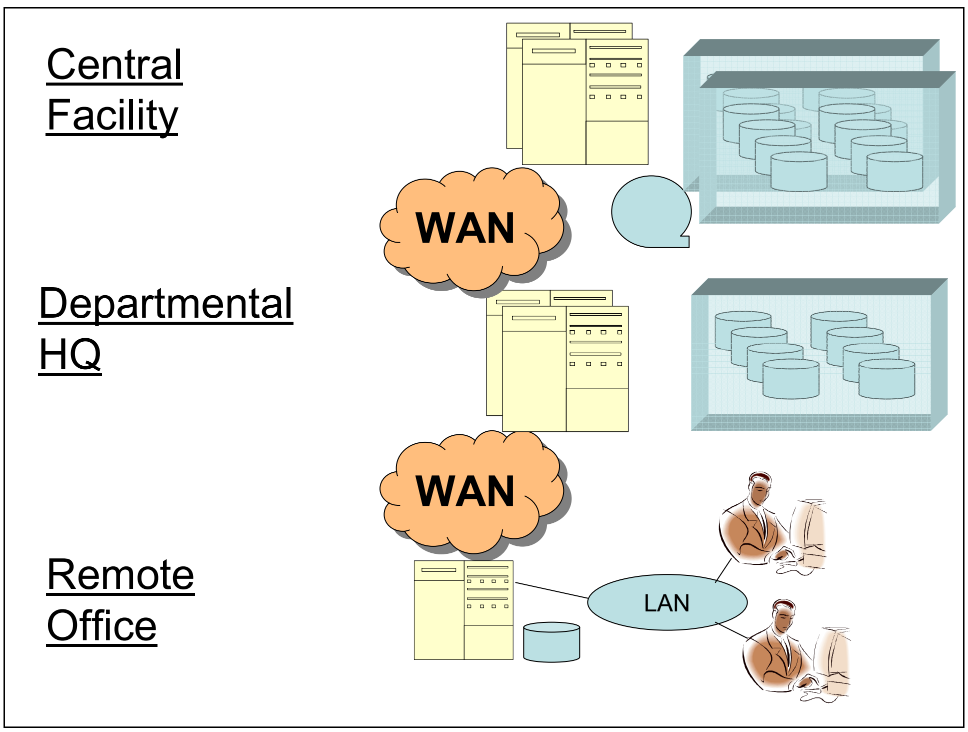

Figure 8-4: Enterprise Content Management Architecture

Implementing this architecture requires the use of additional servers beyond those required in the file-server-only approach. An enterprise content management system typically requires one big content metadata server and content server deployed at a central location, surrounded by 'intelligent caching' servers at multiple logical 'points of presence' around the network.

The document access interface automatically determines which point of presence should be accessed in order to conduct queries and retrieve documents. The preferred 'point of presence' may change from moment to moment and query to query, depending on the conditions the content management system finds at run time.

***Key Observation*:** We saw anecdotal evidence of limited content management system deployment in some departmental inventories, but <u>no indication of any sort of common statewide standard for content management.</u>

Goal 2, Objective 3 of the 2006 update to the IT Strategic Plan describes plans for enhancing existing document management systems. The examples cited in the Strategic Plan are typically for high-volume, transactional document imaging and processing applications.

The enterprise content management opportunity described here suggests a much broader potential application of content management, to be applied across the board as a replacement for file-server and local-disk based content storage. Part of the justification for this is file server consolidation. Better document access, security, and management, however, are a MUCH more significant justification.

Benefits:

- **Eliminates the need for local backup operations** – Only the 'copy of record' of all documents stored at the highest level of the hierarchy requires backup.
- **Document search is independent of document content storage.** Search repositories and content repositories can be replicated to address network performance issues, and the ECM system keeps all repositories synchronized.
- **Inherent high-availability** – If the local ECM cache server instance fails, then the next server level up in the hierarchy can provide ECM, transparently. When the local server comes back up, local users see a performance improvement, but they've never lost access to their documents during the outage.
- **Greatly improves document accessibility** – Instead of searching for file names within a convoluted directory structure, or searching for a word or phrase in a full-text repository, users can quickly search for documents based on secondary, user-definable criteria. Content access rights can be managed on a highly granular basis, which greatly improves document-level security.
- **Reduces overall storage requirements** – By effectively eliminating casual duplication of content.
- **Simplifies dispersed server management**. The ECM system manages both repositories and caches as a single, integrated whole. Administrators define policies, which apply to all servers in the system – they do not have to maintain individual server-level policies.

Limitations:

- **Increases demand for network bandwidth and reliability**. If a document is not found locally, and therefore must be retrieved from a higher level in the hierarchy, response time is delayed if the network is congested or unavailable.

Intelligent pre-caching techniques can be used in an ECM solution that is not available to classic file system approaches or WAFS.

- **Increases complexity**. ECM systems are much more powerful than any conventional file system. They can address a range of requirements that are beyond the reach of conventional file servers. Such power comes at a cost, however, in terms of learning curve for new tools and overall complexity. Proper planning and training can address the complexity, but it should be acknowledged.

8.3.3. Cost and Value

This cost and value analysis is a “What if?” scenario looking at file server consolidation, without WAFS or ECM to a maximum of two servers per site.

The inventory collected identified 882 file servers from 78 sites. They differ widely in processor, memory, storage, and operating system specifications. We attempted to find the most common specifications for these features and decided upon Microsoft Windows Server 2003 and Intel 2-processor servers as the proxy for servers. The model assumes that the sites that currently implement SAN storage only have one subsystem and that those that do not currently implement SAN storage will purchase only one subsystem. The following cost and value analysis is meant to be an illustration of the potential savings from consolidating file servers to a maximum of two per site.

Assumptions:

- Analysis based on the inventory sample of 882 physical file servers.
- This scenario does not account for any servers being virtualized.
- For larger-scale sites with more than two file servers, consolidation of all file servers to exactly two physical file servers is practical.
- The old file servers will be discarded and not reused.
- Windows Server 2003 and Intel 2-processor systems will be used as a proxy for the servers being eliminated.
- Per-user licensing costs will remain the same, since the user counts are assumed not to change.
- A SAN-based storage subsystem will be required at each non-remote site that implements file server consolidation. Sites with pre-existing SAN’s would be able to re-use this resource.
- Sites without SAN infrastructure will need to implement it. A reasonable cost estimate for a single iSCSI storage subsystem with sufficient disk and I/O throughput capacity to support file server head-end consolidation is $25,000, which includes required SAN management software.
- Each site is assumed to implement only one SAN subsystem.
- A yearly warranty fee of 20% of the SAN purchase cost begins after the second year.
- Storage administration tasks in a SAN environment are slightly more complex than in the current direct-attached case; increase administration costs by 10% to account for this.
- Project implementation cost of $2M—includes consultant fees, solution design, and implementation. Specific project implementation costs are highly variable

and are subject to numerous factors and design decisions. We supply a judgment cost here to provide a more accurate analysis.

- The FTE[25] to server ratio used for generic servers is 1:15.

Environment

The server configurations priced in the TCO analysis for File Server Consolidation are:

Consolidated file servers:

82 2-way Xeon DP processors, 8 GB of memory, SAN attached, Windows Server 2003 Enterprise Edition

45 2-way Xeon DP processors, 2 GB of memory, SAN attached, Windows Server 2003 Standard Edition

Current file servers:

882 2-way Xeon DP processors, 2 GB of memory, SAN attached, Windows Server 2003 Standard Edition

NOTE: The current file server configuration approximates the most common specifications of file servers reported in the server inventory.

Table 24 shows the major components of the TCO model priced in this study.

	Current File Server Scenario	*Recommended File Server Scenario*
Number of physical servers:	882	127
Number of virtual servers:	0	0
Number of software licenses:	882	127
Number of FTEs:	58.8	8.47
Data center floor space (square feet):	3784.9	544
Project Implementation:	0	$2M

Table 24: Major TCO components

[25] FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance. FTE's by workload are based on CIOView standard metrics.

TCO Analysis Results

Table 25 shows the estimated cost per year of the current file server environment as modeled in this TCO. The storage cost reflects the three current sites that have SAN subsystems attached to file servers.

Current File Server	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	606,000	606,000	606,000	821,000	821,000	3,459,000	9.63%
Software	441,000	441,000	441,000	441,000	441,000	2,205,000	6.14%
Facilities	817,000	817,000	817,000	817,000	817,000	4,087,000	11.38%
Personnel	5,212,000	5,212,000	5,212,000	5,212,000	5,212,000	26,058,000	72.53%
Implementation	-	-	-	-	-	-	0.00%
Storage	15,000	15,000	30,000	30,000	30,000	120,000	0.33%
Total Cost of Ownership	7,091,000	7,091,000	7,106,000	7,321,000	7,321,000	35,930,000	100.00%

Table 25: Estimated total cost of current file server environment

Table 26 shows the estimated cost per year of the consolidated file server solution.

File Server Opportunity	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	95,000	95,000	95,000	123,000	123,000	531,000	5.38%
Software	93,000	93,000	93,000	93,000	93,000	465,000	4.72%
Facilities	150,000	150,000	150,000	150,000	150,000	749,000	7.59%
Personnel	825,000	825,000	825,000	825,000	825,000	4,127,000	41.83%
Implementation	741,000	667,000	667,000	-	-	2,074,000	21.02%
Storage	240,000	240,000	480,000	480,000	480,000	1,920,000	19.46%
Total Cost of Ownership	2,144,000	2,070,000	2,310,000	1,672,000	1,672,000	9,867,000	100.00%

Table 26: Estimated total cost of file server consolidation recommendation

The implementation costs include disposing of 755 old servers, installing 82 new file servers, connecting 96 servers to a SAN subsystem, and a $2M charge for project implementation spread over three years.

The storage cost accounts for 45 non-remote sites that have to buy new SAN subsystems plus three sites that currently have the subsystems attached to file servers, for a total of 48 subsystems.

Personnel costs include a 10% administrative cost increase for the added complexity of managing the SANs.

Table 27 shows the potential raw savings numbers—for each category across five years—as the difference between the cost of the consolidated file server recommendation and the current file server environment. The total potential savings is about $26M over five years.

Value of Opportunity (Cost shows as negative)	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware	511,000	511,000	511,000	698,000	698,000	2,928,000
Software	348,000	348,000	348,000	348,000	348,000	1,740,000
Facilities	667,000	667,000	667,000	667,000	667,000	3,338,000
Personnel	4,387,000	4,387,000	4,387,000	4,387,000	4,387,000	21,931,000
Implementation	(741,000)	(667,000)	(667,000)	-	-	(2,074,000)
Storage	(225,000)	(225,000)	(450,000)	(450,000)	(450,000)	(1,800,000)
Total Savings	4,947,000	5,021,000	4,796,000	5,649,000	5,649,000	26,063,000

Table 27: Value of file server consolidation recommendation

Table 28 depicts the total cost and savings per category.

Costs	Current File Server	File Server Opportunity	Savings	Value	% Savings
Hardware	3,459,000	531,000	Hardware Savings	2,928,000	84.65%
Software	2,205,000	465,000	Software Savings	1,740,000	78.91%
Facilities	4,087,000	749,000	Facilities Savings	3,338,000	81.67%
Personnel	26,058,000	4,127,000	Personnel Savings	21,931,000	84.16%
Implementation	-	2,074,000	Implementation Savings	(2,074,000)	N/A
Storage	120,000	1,920,000	Storage Savings	(1,800,000)	(15.00)
Total Cost	$ 35,930,000	$ 9,867,000	Total Savings	$ 26,063,000	72.54%

Table 28: Total cost and savings per category

The savings in every category but storage are significant, ~80%, because we are reducing the server count by 85%. The savings incurred because of the server reduction are partly offset by the implementation and storage costs of the consolidated file server solution. The total savings is reduced to 72.5%, $26M over five years.

Figure 8-5 compares the total costs per year of each environment.

Summary	Year 1	Year 2	Year 3	Year 4	Year 5	Total	Savings
Current File Server	7,091,000	7,091,000	7,106,000	7,321,000	7,321,000	35,930,000	26,062,000
File Server Opportunity	2,144,000	2,070,000	2,310,000	1,672,000	1,672,000	9,868,000	

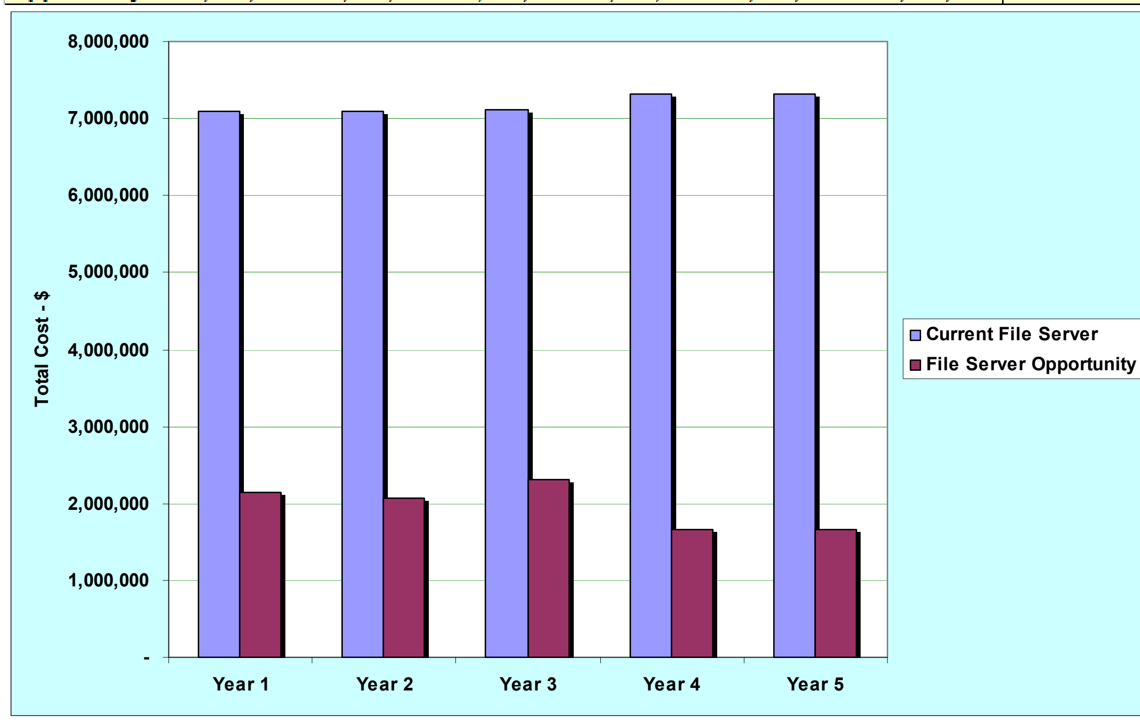

Figure 8-5: File server consolidation recommendation total cost comparison

Key Financial Result:
The estimated total savings from consolidating file servers from 882 to 127 servers is $26M, or 72.5% over five years.

Energy Conservation:

By disposing of 755 physical file servers, the State can conserve 4,417 megawatt-hours per year. This savings translates to an 86% reduction in energy consumption relative to the current file server environment described in this "What if?" scenario. This energy consumption calculation does not include the energy consumed by the new SAN equipment purchased as part of the recommendation.

8.3.4. Risks

There are risks that need to be addressed related to file sharing and content management consolidation. In addition, issues and concerns that are presented in Chapter 10 - Issues and Concerns should be reviewed:

Risk	Mitigation
Network – Relocating file servers can affect network loads.	• If sufficient WAN bandwidth does not exist, it can be acquired. • Network cost to the department is a risk and a concern for some departments. These are offset by improved manageability and reduced server costs. • Metropolitan Area Network (MAN) architecture should be considered or expanded. • As part of the project, a detailed migration plan should be developed.
Complexity – WAFS and ECM have complex infrastructure requirements.	• A detailed project plan must be in place. • The costs would be determined as part of the planning process. • The complexity needs to be evaluated in light of the benefits. • This should be an enterprise-wide offering and not engineered by every department.
Competing Priorities	• The priority has to be compared with other projects. • The long-term benefit needs to be

	appreciated as part of project evaluation.
Resistance to Change	• "File servers already work why change them." The cost to maintain the existing file server infrastructure is too high.
Cost	• The cost is offset by reduced server infrastructure. • Better content discovery and management.

8.4. E-mail

E-mail is run as separate applications within State departments. Each department is responsible for designing, engineering, deploying, and maintaining its e-mail system. DTS has a cross-department e-mail offering that began in 2005 that a few departments currently use. Alternatively, DTS could evaluate an outsourced e-mail offering for the State if security, confidentiality, service levels, and all requirements can be met.

8.4.1. Recommendation

A plan should be developed to convert all departments to a common State e-mail system over a three-year period. Complete the e-mail architecture, engineering, deployment plan, deployment schedule, and pilot in the first year. Convert all departments to the new e-mail system in the second and third years.

8.4.2. Detail

From the inventory sample collected, there were 359 servers across 31 departments classified with a primary function of e-mail. 11 departments had 10 or more e-mail servers and 4 departments had 20 or more e-mail servers. The DTS accounted for 94 e-mail servers. It is assumed that a number of these are being hosted for other departments or organizations.

Over the past decade, e-mail has been rapidly growing in complexity. E-mail systems today must handle a larger volume but also must deal with spam, virus protection, security threats, gateways, anywhere mobile devices, remote access, business continuity, and appropriate use.

Another big issue that has risen is e-discovery. E-mail systems now are subject to discovery for litigation. This is causing changes in how, what, and for how long e-mail messages are subject to archiving and retention. As requirements in this area change, policies need to be adapted or revised.

The requirements and expertise to design, engineer, and maintain e-mail and messaging systems is much greater than five and ten years ago. The requirement for the number of support personnel for each e-mail system has also grown.

State E-mail Architecture:

There are several architecture designs that could be implemented to support a State e-mail system. For example;

- One e-mail system with e-mail servers strategically deployed to remote sites with a large numbers of users; all e-mail systems managed by a central group.
- One e-mail system with all servers centrally located; all users access their e-mail across the WAN.
- Several federated e-mail systems; larger departments, or when requirements dictate, retain their own e-mail system but are managed and integrated by a central group.

We would recommend one e-mail system with e-mail servers strategically deployed to remote sites with large numbers of users.

By using a clustered and replicated system design, a higher availability and improved business continuity could be achieved over what most departments have today. A statewide e-mail system design could still allow for mailbox sizes and retention policies to be set at a group or department level.

DTS could evaluate an outsourced e-mail offering for the State if security, confidentiality, service levels, and all requirements can be met. A careful requirements collection and evaluation process is needed first. E-mail is a critical application at the State and its integration with directories and other applications is growing. E-mail systems are converging into unified communications systems supporting e-mail, instant messaging, mobility, voice messaging, collaboration, search, and content correlation. Outsourced offerings should be evaluated to ensure that they meet these requirements.

8.4.3. Cost and Value

This cost and value analysis is a "What if?" scenario looking at consolidating 250 e-mail servers down to 30 e-mail servers The inventory collected identified 376 e-mail servers, 359 of which are Intel processor-based. They differ widely in processor, memory, storage, and operating system specifications. We attempted to find the most common specifications for these features and decided upon Microsoft Windows Server 2003 and Intel 2-processor servers as the proxy for servers. The scenario also assumes that all servers are SAN attached when in reality they are not. For simplicity, the e-mail software chosen for the TCO model is Exchange Server 2007, although the inventory shows a variety of e-mail software deployed across the departments. The following cost and value analysis is meant to be an illustration of the potential savings from consolidating e-mail servers to a centralized, statewide e-mail system.

Assumptions:

- Analysis based on the inventory sample of 359 e-mail servers.
- 250 of these servers could be consolidated with a State e-mail system.
- The remaining 109 servers serve other functions that would not be consolidated.
- The State e-mail system would support 150,000 mailboxes[26].
- One e-mail server would support 8,000 mailboxes.
- There will be a 4:1 redundancy ratio of servers for high-availability.
- Four additional servers are required for support services.
- Two servers with virtual machines will be used for development and testing.
- External storage for mailboxes would be the same and is not factored in.
- Project implementation cost of $5M—includes consultant fees, solution design, and implementation. Specific project implementation costs are highly variable and are subject to numerous factors and design decisions. We supply a judgment cost here to provide a more accurate analysis.
- The FTE[27] to server ratio used for e-mail servers is 1:7.

[26] This is an assumption only. We do not have data to indicate how many mailboxes would have to be support.

Environment

The server configurations priced in the TCO analysis for the E-mail Consolidation scenario are as follows:

Consolidated e-mail servers:

> 24 4-way Xeon MP processors, 32 GB of memory, SAN attached, Exchange Server 2007 Enterprise Edition
>
> 6 2-way Xeon DP processors, 4 GB of memory, SAN attached, Exchange Server 2007 Enterprise Edition

NOTE: 19 Xeon MP servers are used to host the 150,000 mailboxes stated in the assumptions with five extra Xeon MP servers used for redundancy. Four Xeon DP servers are used for support services and two Xeon DP servers host the virtual machines for development and testing.

Current e-mail servers:

> 250 2-way Xeon DP processors, 4 GB of memory, SAN attached, Exchange Server 2007 Standard Edition

NOTE: The server configuration used to represent all current e-mail servers approximates the most common specifications of the servers used for e-mail as reported in the server inventory.

Table 29 shows the major components of the TCO model priced in this study.

	Current E-mail Model	*Recommended E-mail Model*
Number of physical servers:	250	30
Number of virtual servers:	0	8
Number of software licenses:	250	38
Number of FTEs:	36	5
Data center floor space (square feet):	1419.2	136.4
Project Implementation:	0	$5M

Table 29: Major TCO components

[27] FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance. FTE's by workload are based on CIOView standard metrics.

TCO Analysis Results

Table 30 shows the estimated cost per year of the current e-mail solution as modeled in this TCO.

Current e-mail	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	193,000	193,000	193,000	254,000	254,000	1,086,000	5.67%
Software	160,000	160,000	160,000	160,000	160,000	800,000	4.18%
Facilities	286,000	286,000	286,000	286,000	286,000	1,430,000	7.47%
Personnel	3,165,000	3,165,000	3,165,000	3,165,000	3,165,000	15,827,000	82.68%
Implementation	-	-	-	-	-	-	0.00%
Total Cost of Ownership	3,804,000	3,804,000	3,804,000	3,865,000	3,865,000	19,142,000	100.00%

Table 30: Estimated total cost of current e-mail environment

Table 31 shows the estimated cost per year of the consolidated e-mail solution. The Implementation costs include new server provisioning, old server disposal, and, most importantly, a $5M charge for project implementation spread over three years.

e-mail Opportunity	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	92,000	92,000	92,000	103,000	103,000	480,000	5.84%
Software	45,000	45,000	45,000	45,000	45,000	226,000	2.76%
Facilities	36,000	36,000	36,000	36,000	36,000	178,000	2.17%
Personnel	461,000	461,000	461,000	461,000	461,000	2,304,000	28.07%
Implementation	1,689,000	1,667,000	1,667,000	-	-	5,023,000	61.17%
Total Cost of Ownership	2,323,000	2,300,000	2,300,000	644,000	644,000	8,211,000	100.00%

Table 31: Estimated total cost of e-mail recommendation

Table 32 shows the potential raw savings numbers—for each category across five years—as the difference between the cost of the e-mail consolidation recommendation and the current e-mail environment. The total potential savings is about $11M over five years.

Value of Opportunity (Cost shows as negative)	Year 1	Year 2	Year 3	Year 4	Year 5	Total Savings
Hardware	101,000	101,000	101,000	151,000	151,000	606,000
Software	115,000	115,000	115,000	115,000	115,000	574,000
Facilities	250,000	250,000	250,000	250,000	250,000	1,252,000
Personnel	2,704,000	2,704,000	2,704,000	2,704,000	2,704,000	13,523,000
Implementation	(1,689,000)	(1,667,000)	(1,667,000)	-	-	(5,023,000)
Total Savings	1,481,000	1,504,000	1,504,000	3,221,000	3,221,000	10,931,000

Table 32: Value of E-mail Recommendation

Table 33 depicts the total cost and savings per category.

Costs	Current e-mail	e-mail Opportunity	Savings	Value	% Savings
Hardware	1,086,000	480,000	Hardware Savings	606,000	55.80%
Software	800,000	226,000	Software Savings	574,000	71.75%
Facilities	1,430,000	178,000	Facilities Savings	1,252,000	87.55%
Personnel	15,827,000	2,304,000	Personnel Savings	13,523,000	85.44%
Implementation	-	5,023,000	Implementation Savings	(5,023,000)	N/A
Total Cost	$ 19,142,000	$8,211,000	Total Savings	$ 10,931,000	57.10%

Table 33: Total cost and savings per category

The savings across the board are due to the replacement of old e-mail servers with fewer, newer servers with lower power requirements. The following factors contribute to the cost benefits:

- lower power requirements per server result in lower data center utility bills
- lower capital server costs
- fewer software licenses
- less data center floor space utilization
- less overall cabling and networking
- less maintenance and operating costs
- fewer support personnel

Even with the one-time charge of $5M for project implementation, the State could potentially save $11M over five years, or 57% of its estimated current e-mail environment costs.

Figure 8-6 compares the total costs per year of each environment. The costs of the e-mail consolidation recommendation decline drastically after the third year because it is assumed that the implementation is paid for by the third year.

Summary	Year 1	Year 2	Year 3	Year 4	Year 5	Total	Savings
Current e-mail	3,804,000	3,804,000	3,804,000	3,865,000	3,865,000	19,142,000	10,931,000
e-mail Opportunity	2,323,000	2,300,000	2,300,000	644,000	644,000	8,211,000	

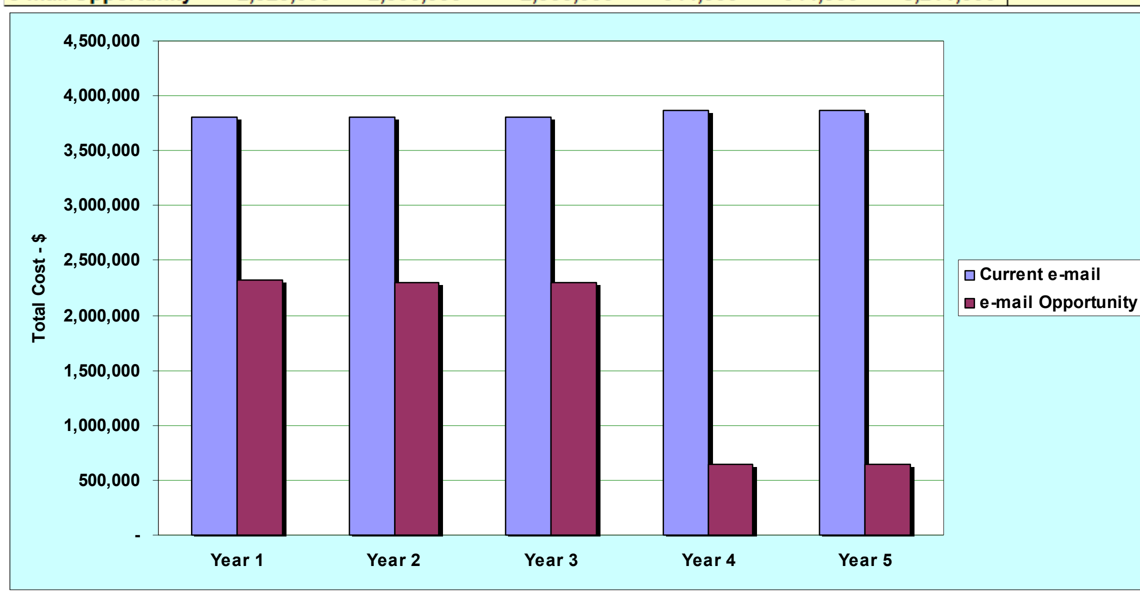

Figure 8-6: E-mail recommendation total cost comparison

Key Financial Result:
The State could potentially save $11M, or 57% over five years, of its estimated current e-mail environment costs, by consolidating 250 e-mail servers down to 30 e-mail servers.

Energy Conservation:

By disposing of 226 e-mail servers, the State can conserve about 1,461 megawatt-hours per year. This savings translates to a 90% reduction in energy consumption relative to the current e-mail environment depicted in this "What if?" scenario.

Other benefits:

- An engineering team dedicated to e-mail could provide better quality services.
- A higher level of security can be maintained.
- Response to virus and security threats can be quicker without every department having to have resources responding.
- Standards can be maintained.
- High availability solutions can be engineered.
- Continuity of Government e-mail strategies are simpler to engineer and maintain.
- Better support for e-discovery.

8.4.4. Risks

There are risks that need to be addressed related to e-mail consolidation. In addition, issues and concerns that are presented in Chapter 10 - Issues and Concerns should be reviewed:

Risk	Mitigation
Network – Relocating e-mail can affect network loads.	• If sufficient WAN bandwidth does not exist, it can be acquired. • Network cost to the department is a risk and a concern for some departments. These are offset by improved e-mail management, reduced servers, and improved benefits stated above. • Metropolitan Area Network (MAN) architecture should be considered or expanded. • As part of the project, a detailed migration plan should be developed.
Directory – departments do not share the same Active Directory.	• How directories are kept in sync and trust relationships need to be part of the engineered solution.
Complexity – All State personnel would have to have access to the consolidated e-mail.	• This would have to be considered as part of the engineering design. • Actually, a statewide e-mail system once implemented would be far less complex than every department

	engineering and managing their own e-mail systems.
Competing Priorities	• The priority has to be compared with other projects. • The long-term benefit needs to be appreciated as part of project evaluation.
Resistance to Change	• "Our e-mail system works just fine!" The cost and complexity to maintain the existing e-mail infrastructure is too high. • Continuity of Government and high-availability need to be engineered into existing department e-mail systems. • E-mail systems have security risks. • One team devoted to e-mail can provide better service levels than every department engineering their own solution. • Better support for e-discovery.
Cost	• The cost is offset by reduced server infrastructure. • Consolidating engineering resources would produce better quality at reduced cost.

8.5. *Virtualization*

Often, virtualization is the first thought that comes to mind when server consolidation is discussed. In fact, it does present the largest opportunity for reducing physical servers. However, virtualization does not eliminate most software licensing costs or operating system and application support costs. We recommend looking at the other opportunities for server consolidation first as they have the potential for reducing software licensing, operating system support, and application support costs.

8.5.1. Recommendation

The State CIO should set a goal to convert 20% of the existing servers to virtual machines over a three-year period. In addition, a plan should be developed with the cooperation of the departments to meet this goal. A ***simple*** quarterly tracking spreadsheet/system should be set up to record information by department. The spreadsheet and progress should be reported quarterly to the State CIO.

8.5.2. Detail

Several departments have begun or are planning to utilize virtual machines as a way to reduce physical machines. We recommend that this effort be acknowledged and tracked.

A goal to convert 20% of the exiting servers to virtual machines over a three-year period should be manageable. Other companies have achieved goals of 25-35% and some achieve 50%.

Deploying virtual machines can also provide added advantages such as:[28]

- Backup – the ability to backup entire virtual machines. If there is limited application data and it is included with the virtual machine, it can be backed up with the system as an intact image.
- Load Balance – virtual machines can be moved between physical machines if loads need to be balanced or additional capacity is needed at peak times.
- DR and Business Continuity – entire images of virtual machines can be used for recovery for DR and Business Continuity.
- Development and Testing – virtual machines can provide snapshots and rollbacks for repeating test cases and multiple tiered environments.

Regardless of utilization, not all servers are good candidates for virtualization. Some systems have special hardware or systems drivers that are not compatible with virtualization. In addition, occasionally, applications can have an issue running in a virtual environment; this is usually caused by direct access to hardware.

[28] IDC provides a good research paper, ten pages, titled “Industry Development and Models: IDC’s Server Virtualization Maturity Index,” December 2006, IDC #204893. This paper provides IDC’s opinion on server virtualization maturity today and its future outlook

Two issues that frequently arise with the use of virtual machines are:

- Over commitment of system resources – the first resource that often is over committed is real memory. There needs to be sufficient available real memory on the host computer to support the virtual machines. The other resources that can be over committed are network and disk I/O bandwidth. Performance monitoring of the host computer should continue to be performed to make sure resources are not over committed.
- Proliferation of virtual machines – since virtual machines are easy to create and do not require hardware purchase, new ones may be deployed at a faster rate than if physical machines were used. Virtual machines still require software licensing, operating system support, and application support. Uncontrolled growth of virtual machines adds to support staff workloads—usually the highest factor in TCO.

8.5.3. Cost and Value

This cost and value analysis is a “What if?” scenario looking at a 20% conversion of existing physical machines to virtual machines.

The inventory identified 6,000+ servers. The total number is unknown, but we assume there are over 9,000 servers. We will also use Microsoft Windows and Intel 2-processor servers as the proxy for servers. There is also good opportunity for UNIX servers to be converted to virtual machines. This scenario also does not phase in the conversion of servers over time but is intended to show the overall cost reduction for moving physical machines to virtual machines; the actual cost reduction would be realized across several years.

Assumptions:

- There is an installed base of 9,000 servers.
- 20% of this installed base—1,800 physical machines—will be converted to virtual machines on new hardware.
- Old hardware—1800 servers—will be retired.
- Windows Server 2003 and Intel 2-processor systems will be used as a proxy for the servers being converted.
- The virtual to physical server ratio will be 8:1 (4 per processor).
- Each virtual machine will have 1GB dedicated real memory.
- Physical servers will be connected to SAN and virtual machine images will be stored on SAN.
- No new software licensing is required on the virtual machines.
- New virtualization software is required on the host physical machines.
- The conversion of servers is all at once, not spread across three years.
- Project implementation cost of $1.5M—includes consultant fees, solution design, and implementation. Specific project implementation costs are highly variable, are subject to numerous factors, and design decisions. We supply a judgment cost here to provide a more accurate analysis.

- The FTE[29] to server ratio used for generic servers is 1:15.

Environment

The server configurations priced in the TCO analysis for the Virtualization scenario are:

Physical hosts for virtualized servers:

2-way Xeon Core 2 Duo-based processors, 8 GB of memory, SAN attached, VMware* ESX

NOTE: The server configuration approximates the most common specifications of servers reported in the server inventory plus six more GB of memory for the virtual machines. VMware ESX is the virtual server software chosen for this TCO model.

Current non-virtualized servers:

2-way Xeon DP processors, 2 GB of memory, SAN attached

NOTE: The server configuration approximates the most common specifications of servers reported in the server inventory.

Table 19 shows the major components of the TCO model priced in this study.

	Current In-Department Model	*Recommended In-Department Model*
Number of physical servers:	1800	225
Number of virtual servers:	0	1800
Number of software licenses:	1800	2025
Number of FTEs:	120	111
Data center floor space (square feet):	7724.3	965.5
Project Implementation:	0	$1.5M

Table 34: Major TCO components

[29] FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance. FTE's by workload are based on CIOView standard metrics.

TCO Analysis Results

Table 35 shows the estimated cost per year of the current non-virtualized server environment as modeled in this TCO.

Current No Virtualization	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	1,237,000	1,237,000	1,237,000	1,675,000	1,675,000	7,060,000	9.37%
Software	900,000	900,000	900,000	900,000	900,000	4,500,000	5.97%
Facilities	2,123,000	2,123,000	2,123,000	2,123,000	2,123,000	10,617,000	14.09%
Personnel	10,636,000	10,636,000	10,636,000	10,636,000	10,636,000	53,180,000	70.57%
Implementation	-	-	-	-	-	-	0.00%
Total Cost of Ownership	14,896,000	14,896,000	14,896,000	15,334,000	15,334,000	75,357,000	100.00%

Table 35: Estimated cost of non-virtualized server environment

Table 36 shows the estimated cost per year of the virtualized server solution. The implementation costs include provisioning 225 new servers, disposing of 1800 old servers, and a $1.5M charge for project implementation spread over three years.

Virtualization Opportunity	Year 1	Year 2	Year 3	Year 4	Year 5	Total	% of TCO
Hardware	180,000	180,000	180,000	228,000	228,000	995,000	1.62%
Software	2,474,000	1,452,000	1,452,000	1,452,000	1,452,000	8,284,000	13.47%
Facilities	265,000	265,000	265,000	265,000	265,000	1,327,000	2.16%
Personnel	9,838,000	9,838,000	9,838,000	9,838,000	9,838,000	49,192,000	80.00%
Implementation	688,000	500,000	500,000	-	-	1,688,000	2.75%
Total Cost of Ownership	13,446,000	12,236,000	12,236,000	11,784,000	11,784,000	61,486,000	100.00%

Table 36: Estimated cost of virtualization recommendation

Table 37 shows the potential raw savings numbers—for each category across five years—as the difference between the cost of the virtualization recommendation and the current non-virtualized environment. The total potential savings is about $14M over five years.

Value of Opportunity (Cost shows as negative)	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware	1,057,000	1,057,000	1,057,000	1,447,000	1,447,000	6,065,000.00
Software	(1,574,000)	(552,000)	(552,000)	(552,000)	(552,000)	(3,784,000.00)
Facilities	1,858,000	1,858,000	1,858,000	1,858,000	1,858,000	9,290,000.00
Personnel	798,000	798,000	798,000	798,000	798,000	3,988,000.00
Implementation	(688,000)	(500,000)	(500,000)	-	-	(1,688,000.00)
Total Savings	1,450,000	2,660,000	2,660,000	3,550,000	3,550,000	13,871,000.00

Table 37: Value of virtualization recommendation

Table 38 depicts the total cost and savings per category.

Costs	Current No Virtualization	Virtualization Opportunity	Savings	Value	% Savings
Hardware	7,060,000	995,000	Hardware Savings	6,065,000	85.91%
Software	4,500,000	8,284,000	Software Savings	(3,784,000)	-84.09%
Facilities	10,617,000	1,327,000	Facilities Savings	9,290,000	87.50%
Personnel	53,180,000	49,192,000	Personnel Savings	3,988,000	7.50%
Implementation	-	1,688,000	Implementation Savings	(1,688,000)	N/A
Total Cost	$ 75,357,000	$ 61,486,000	Total Savings	$ 13,871,000	18.41%

Table 38: Total cost and savings per category

The savings numbers in Table 38 reflect some of the limitations of a purely virtualized environment. While hardware savings are significant—86% saved in hardware costs—software costs actually increase. The reason that software costs increase is because instead of reducing or even maintaining the same number of server images to manage, the number of server images increases. The number of server images consists of the virtual images plus the images of the host machines—225 in this case.

Savings in hardware support personnel are largely offset by an increase in personnel needed to support the extra images. Table 39 illustrates this point assuming 1800 physical servers are virtualized to 1800 virtual machines on 225 physical servers:

	Before Virtualization	After Virtualization
Hardware FTEs	24	3
Software FTEs	96	108
Total	120	111

Table 39: Comparison of FTEs needed to support a non-virtualized versus a virtualized environment

Figure 8-7 compares the total costs per year of each environment.

Summary	Year 1	Year 2	Year 3	Year 4	Year 5	Total	Savings
Current No Virtualization	14,896,000	14,896,000	14,896,000	15,334,000	15,334,000	75,356,000	13,870,000
Virtualization Opportunity	13,446,000	12,236,000	12,236,000	11,784,000	11,784,000	61,486,000	

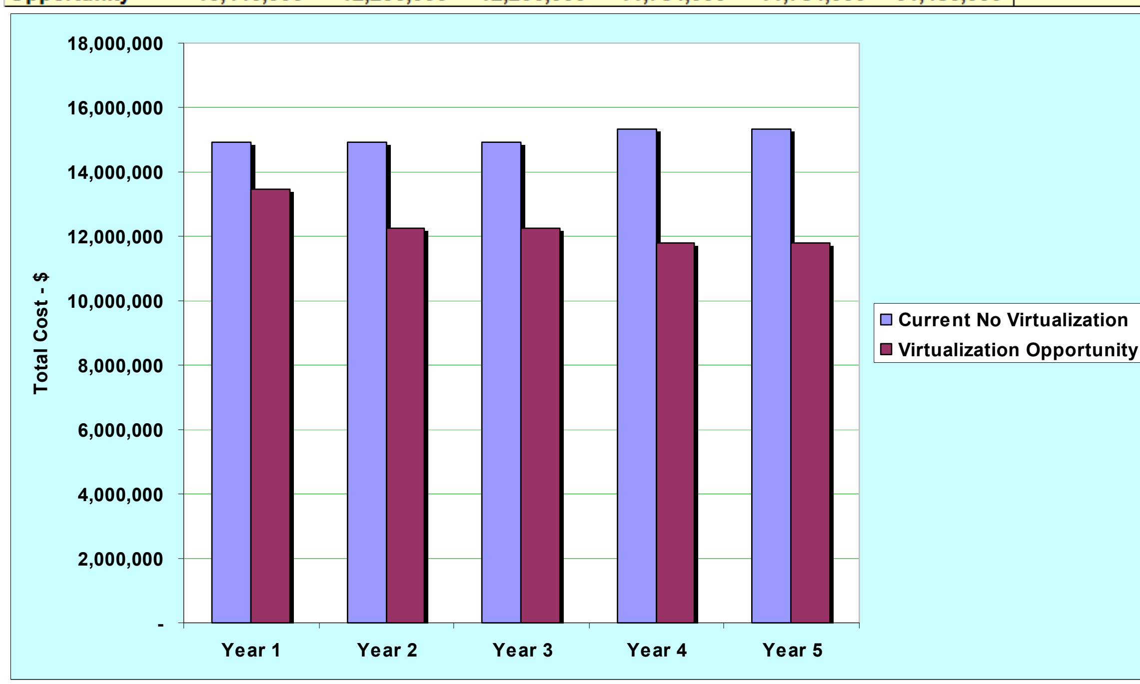

Figure 8-7: Virtualization recommendation total cost comparison

Key Financial Result:
By virtualizing 1800 servers on 225 new servers the State can potentially save $14M, 18.4% over five years of the cost to maintain the 1800 old physical servers.

Energy Conservation:

By disposing of 1575 physical servers, the State can conserve about 9,215 megawatt-hours per year. This savings translates to an 88% reduction in energy consumption relative to the current non-virtualized server environment depicted in this “What if?” scenario.

Other benefits:

- Backup and recovery could be simplified.
- Hardware upgrades could occur without affecting server operating system.
- Ability to load-balance.
- Provide for simpler DR and Continuity of Government planning and maintenance.

8.5.4. Risks

There are risks that need to be addressed related to server virtualization. Also, issues and concerns that are presented in Chapter 10 - Issues and Concerns should be reviewed:

Risk	Mitigation
Over commitment of physical system resources	• Capacity and performance planning has to be done to ensure physical servers resources are not over committed.
Proliferation of virtual machines	• Due diligence must be maintained to not deploy additional servers. Software licenses must be maintained for each virtual server. • Software support personnel are still required to support virtual servers.
Complexity	• Maintaining a virtual server environment does require some planning and training. • The training required is usually 3-5 days for two primary support people and 1 day for other support personnel.

9. Additional Opportunities

Additional opportunities are consolidation opportunities that should be considered by the State or departments. There is some evidence in the server inventory to support these opportunities but further evaluation is needed to qualify the opportunity.

This section should be viewed as suggestions to review and not as recommendations. Departments should use this list of opportunities as part of their own in-department consolidation review to determine if the opportunity should be explored further.

9.1. Database

Opportunity: Departments should review their database servers to determine if there is opportunity to consolidate to fewer servers. Microsoft SQL and Oracle support multiple instances of the database engine on a server. This can allow building high-availability database clusters and sharing it with several databases.

DBMS licenses tend to be relatively expensive and a major portion of TCO on database servers. Consolidating database servers allows for an overall savings in DBMS licensing costs. In addition, high-available server clusters can be design and used to support several database instances.

It is tempting to use virtualization for consolidating underutilized and proliferated database servers. However, DBMS license costs are high and virtualization may not reduce this cost. While it costs more in terms of effort to consolidate database instances onto a smaller number of more-powerful database servers, the payback can be much larger in terms of total operational cost savings.

9.2. Directory Servers

Opportunity:

- Conduct an in-depth analysis of the possibility of defining a common statewide directory services and network infrastructure architecture.
- If a statewide directory service is not feasible, small and medium size departments should consider participating in a consolidated directory service managed by DTS.

There were 430 servers in the inventory identified as performing some kind of directory service[30].

[30] E.g. – LDAP, DNS, DHCP, Active Directory, Netware Directory Services, etc. – Any service which accepts a request for the name or address of another computer, or provides a dynamic address assignment for a new addition to the network was considered a 'Directory' server.

Two things are clear about Directory Servers:

1. There is little consistency across departments in terms of how directories are operated.
2. If the State resembles most other organizations, then the majority of the servers used to implement directory-management schemes across the State are likely to be underutilized.

The diversity of directory solutions in use creates unnecessary cost and complexity. However, the challenges involved in delivering a statewide solution are very daunting. Still, a study should be done to determine if there are options for consolidating directories or other alternatives for building a federated directory structure.

DTS could offer and manage a shared directory service for use by departments. This service would require DTS to manage directory servers at remote locations closer to users and member servers.

9.3. Web Servers

Opportunity: Look at combining websites from multiple web servers on a single web farm.

Many data-center-level sites inventoried included multiple servers that are dedicated to web functions. Some of these servers were noted as serving internal clients only; others are publicly viewable. The exact breakdown of Internet vs. Intranet servers cannot be determined from the information available.

Many applications deploy their own web frontend on dedicated servers. These can be underutilized servers serving one application or function. These websites can be combined on a shared web farm that can provide high-availability and load balancing.

9.4. Remote Access Servers

Opportunity: Evaluate the utilization of Terminal Service servers. Move servers into server farms when possible to share resources and reduce licensing costs; maximize the number of users per server and run servers as virtual machines. It may be possible to run four to eight Terminal Service servers on a two-processor physical machine.

The use of Microsoft Terminal Service and Citrix is a good option for reducing network requirements for network intensive client server applications. It is also used to reduce the application management overhead and licensing cost for users with limited application needs. In some cases, the need for Terminal Service is being replaced as client server applications convert to web front ends.

Fourteen departments are using Microsoft Terminal Service (mostly with Citrix) on 90 servers. This number may be low, as some Terminal Service servers could have been classified as application servers instead.

There could be opportunity for hardware and software licensing cost reductions by combining servers into Citrix Terminal Service farms and combining the workloads and licensing costs. If consolidation of these servers occurred at a DTS data center, there could be an opportunity to reduce the number of Terminal Service servers, possibly in half. However, the architecture of the application back-ends and the Active Directory domains that the servers are a part of would be a factor and would have to be evaluated. There is also the opportunity for consolidating Terminal Service servers locally within a department.

There are Windows resource limitations when consolidating Terminal Servers. Often, the processing power of the hardware is underutilized when the maximum users on the server OS is reached. To overcome this limitation, several Terminal Service servers can be run as virtual machines on one physical server.

9.5. Business Applications

Opportunity:

- Consider using shared landing—combining applications on one system—for the deployment of new applications. If shared landing is not appropriate, consider using a virtual machine.
- For older application servers, at server refresh or other upgrade, consider converting application server to a virtual machine.

There were 1568 servers in the inventory identified as application servers. This is by far the largest single category of servers.

Any application servers that are significantly underutilized may be good candidates for consolidation using virtualization. Virtualizing these servers will not save on OS, DBMS or application licensing costs, but it can save significant amounts of floor space, HVAC, and electrical costs by reducing the absolute number of physical servers in use. The vast majority of both commercial and in-house applications can safely be run under virtualization, though vendor-support challenges can sometimes arise in a virtualized setting.

9.6. Development Systems

Opportunity: Departments should consider using virtual machines for their lab environment. Caution though, this has lead to the proliferation of virtual machines. Each virtual machine still needs to be licensed, patched, and managed,

There were 709 servers identified as development servers. There were 67 development servers running as virtual machines.

The use of virtual machines for development and testing has become a best practice. Not only is there the advantage of reducing physical servers but virtual machines also have many added features for a development and test environment, such as snapshots and roll-backs for testing, quick deployment of new development or test systems, use of multiple images for unit testing, function testing, day-one testing, and production bug fix.

9.7. *Data Backup and Recovery*

Opportunity: Departments should consider undertaking an effort to augment site-level backup and recovery operations with centralized backup and recovery tools. DTS could offer a remote backup and recovery service, possibly deploying backup and recovery servers at large sites.

There were 123 Backup servers in the inventory. Therefore, the potential for direct cost savings via consolidation or elimination of backup servers is not large in comparison to other opportunities identified.

However, a cursory review of the backup architecture used by the vast majority of sites shows two things:

1. Backups are typically done over the network, and
2. Backups are typically performed to local (on-site) tape or other archival media storage devices.

This means that every department and site is responsible for developing and maintaining its own backup and recovery approach, testing it periodically, arranging for secure offsite storage, etc. – all of the details that go into implementing a safe and secure backup strategy.

As the news has reported many times of late, failure to implement a safe and secure backup strategy can result in embarrassing and potentially very costly confidential data disclosures.

Technology is now available that could allow the backup infrastructure of many sites and departments to be augmented by one that allows centralization of many of the detailed steps involved in backup. Using this technology, primary storage would remain local to the site, but backups would be done centrally (either on a statewide basis or at the metropolitan level).

10. Issues and Concerns

While server consolidation provides a good opportunity to reduce costs and make the State's IT server infrastructure more efficient, server consolidation presents issues and concerns. As the project team performed the data collection, interacted with State staff, and interviewed staff and vendors, several issues and concerns were expressed or noted.

The following list of issues and concerns do not represent concerns by all parties or departments but they are significant enough to warrant consideration. We do not provide recommendations on how these issues and concerns should be resolved. However, we agree, they should be addressed as part of any action or project.

10.1. Centralization

One of the big questions raised by commissioning this study was "to what degree is server consolidation within the State intended to centralize IT services and servers?"

This study does suggest that consolidating some IT services and centralizing servers can provide cost savings and improve efficiencies. However, this also assumes that the central infrastructure, IT staff, and architecture are adequate to implement and manage these services, with the same or improved service levels.

Some departments expressed the concern that DTS could not provide the same or improved service levels that the departments themselves provide today. This study does not evaluate the current capability of DTS to provide these services. DTS would have to assure that service levels are met.

Concerns regarding DTS' pricing were also expressed. There is an impression by some departments that DTS' rates are too high and that some individual departments can design less expensive solutions themselves. This study does not evaluate or address DTS' rates for service or its current architecture. It is difficult to compare departments' costs with DTS' costs for similar services. In the case of DTS services, items such as improved security, built-in refresh, and business continuity are included. In the case of the departments, it is likely not all costs of the services are taken into account. This makes it difficult to compare "apples-to-apples." However, DTS should evaluate their rates and cost allocation model and simplify where possible.

The overhead of doing business with DTS or any cross-department organization is perceived as adding additional tasks, causing undue delays, necessitating learning more processes, additional levels of change control, setting differing priorities, and slowing down projects—overall, decreasing agility. It is true that centralizing services does present additional overhead and may decrease agility. DTS needs to evaluate their processes and customers interaction to minimize this impact and be aware of the perception that they may be generating additional overhead to departments ("walk-a-mile in their shoes").

10.2. Department Autonomy

Some concerns about departmental autonomy are covered in the previous Centralization topic. The concern is that DTS does not share the department's business perspective. An autonomous IT department can have more control over prioritizing responses to problems or its own business objectives. In addition, an autonomous IT department might monitor and control costs better.

10.3. Network

If, in the process of consolidation, servers or services are centralized, then the network impact of such a consolidation needs to be addressed. The main concerns of centralizing servers or service, and the network, are:

- **CSGNET backbone capacity** – if servers or services are centralized, this could create a significant increase in backbone utilization. The increased utilization would need to be accounted for prior to the relocation to assure service levels are not impacted. This should be proactive, not reactive.
- **Network high-availability** – if servers or services are moved, then other portions of the WAN topology can become a single point of failure and cause end-user downtime. Network connectivity to main offices need to be engineered for high-availability. Remote offices without high-availability designs need to have contingency plans for business continuity.
- **Endpoint network architecture –** today, end users may not be directly connected to the CSGNET backbone. Instead, offices and users may be connected to a department's own WAN network, traverse other State networks, VPN across the Internet, or a combination of these. Moving servers or services could affect performance of applications for downstream users.
- **Cost** – redesigning WAN connectivity to departments and remote offices, engineering for high-availability, and increase in bandwidth will increase network costs. These increased network costs can offset any cost savings gained by centralized server consolidation.

10.4. DTS Expertise

To run centralized IT services, DTS must maintain adequate expertise to provide the services. This includes:

- **Account management** – interface with agency and departments to understand business needs, requirements and priorities, and effectively convey them to DTS management and staff; reduce overhead impact to departments.
- **Solution architecture and engineering** – develop IT services that are cost effective and meet departmental requirements as well as governmental requirements, such as security, confidentiality, continuity of government, and business continuity. Services need to be designed for high availability all the way to the user.
- **Operations** – the operational staff needs to have the expertise to monitor and maintain services. This includes early detection and quick resolution of problems and issues.
- **Service administration** – ability to administer services quickly with little overhead. For example, e-mail; being able to add, modify, recover, and suspend accounts.

10.5. State IT Staff Constraints

The State is facing attrition in its IT groups. This is due to workforce retirement and state workers leaving for private sector jobs. If servers or services are centralized or moved across departments, this will cause further skill set issues within departments. The State needs a skill set audit, training opportunities, and a plan for ensuring people with the right skills are in the right positions.

10.6. Project Cost / Competing Priorities

Server consolidation has the potential to reduce costs and improve efficiencies long term. However, there are increased costs, required resources, and staff requirements in the short term that are an impact to departments' budgets. Funding for these project costs, resources, and staff is a budget concern. There are other priorities within departments that compete for IT resources and staff time. Server consolidation opportunities need to be prioritized with other departmental projects.

10.7. Server Funding

Many servers were funded specifically by programs, projects, or budget line items. Some of this type of funding is specific and does not provide for consolidated systems. Providers of the funding may claim that funding was for a specific purpose and specific hardware and "their servers should not be consolidated."

11. Appendices

11.1. Appendix A: Departments Participating in Study

The project team would like to thank all those who contributed data for this study. We recognize that this required time and resources to be made available on a short time frame. Thank you for your assistance and support!

Department of Parks and Recreation
Department of Managed Health Care
CA Law Revision Commission
Victim Compensation and Government Claims Board
Department of Corporations
Commission on Aging
Commission on Teacher Credentialing
Department of Financial Institutions
Dept of Food and Agriculture
CA Managed Risk Medical Insurance Board
Board of Chiropractic Examiners
California State Library
Department of Boating & Waterways
Environmental Protection Agency, CA - Responding for Ca Integrated Waste Management Brd, too.
CA Coastal Commission
State Controller's Office
Department of Industrial Relations
Department of Child Support Services
State Board of Equalization
Department of Pesticide Regulation
State Council on Developmental Disabilities
Office of Real Estate Appraisers
CA Public Utilities Commission
CA Conservation Corps
San Joaquin River Conservancy
Office of Environmental Health Hazard Assessment
Department of Forestry & Fire Protection
Department of Alcohol & Drug Programs

Department of Personnel Administration
State Personnel Board
Governor's Office of Executive Information Services

Franchise Tax Board
HIPAA Implementation, California Office of
Employment Development Department

Integrated Waste Management Board
Tahoe Conservancy
Lottery Commission, CA State
Department of Community Services and Development
Department of Technology Services
Department of Water Resources
Department of Motor Vehicles
Department of Health Services
Department of General Services
Department of Transportation
Department of Corrections and Rehabilitation
California Department of Highway Patrol
State Lands Commission

CA Energy Resources, Conservation & Development Commission
Consumer Affairs
Department of Housing & Community Development
Unemployment Insurance Appeals Board
African American Museum
Ca Student Aid Commission
Public Employment Relations Board
Department of Insurance
Department of Conservation
Department of Finance
Department of Social Services
Department of Developmental Services

Department of Aging
Delta Protection Commission
Emergency Medical Services Authority
ISAWS System Support   Office of Systems Integration
CALFED Bay-Delta Program

Department of Alcoholic Beverage Control
Commission on Peace Officers Standards & Training
Fair Employment & Housing, Department of
State Teachers' Retirement System

Department of Fish & Game
Office of Emergency Services
CA Arts Council
California Energy Commission
Department of Mental Health

11.2. Appendix B: Data Acquisition and Processing Approach

130 agencies/departments were invited to participate. We received inventories from a total of 52 (40%) different departments (plus multiple sites from some departments). We received online surveys from a total of 36 departments (28%). Site-level surveys were received from 48 sites.

Departments were able to provide inventories through two different mechanisms:

1. By populating a standardized 'template' spreadsheet using their own tools, techniques, and existing information, with the requested data elements to describe the physical and functional attributes of their servers, on a per site basis. (See Server Information Spreadsheet in the appendix to view the template.)
2. By using Ecora, an automated discovery tool provided by the consultant team, to automatically scan their networks and discover the details of all servers found on those networks.

Both methods worked reasonably well, but there are some limitations/issues that should be kept in mind when reviewing the data.

1. 'Manual' inventories (those not generated by Ecora in this context) proved to be highly variable in terms of content and quality. A few sites provided comprehensive and obviously accurate details concerning their servers. Many sites did not.
2. The Ecora scans were very effective at accurately identifying all of the physical attributes of all of the Windows, Linux, Netware, proprietary UNIX (e.g. – RISC) and VMware servers discovered by the tool. Ecora also identified all of the applications that were installed (via a commercial installation utility) on each server. There was no way, however, for Ecora to automatically identify the 'primary purpose' of each server, which is a key attribute for the consolidation analysis.

We took all of these factors into account in the process of standardizing and loading the various discovery files into the common 'Simple CMDB (configuration management database)' that we created for this project. (Using MS-Access on the SQL Server).

In processing the manual inventories, we followed certain rules in standardizing the data. Respondents were typically very 'free form' in describing the primary function of the servers. The 'standard' definitions suggested in the template were often not followed. We went through an interpretive process (one consistent with the original categorization instructions provided to respondents) in order to populate the primary function column for manual inventories:

- Servers identified as directly supporting all or part of an application (as opposed to providing generic infrastructure services on behalf of multiple applications) were identified as 'Application'
- Servers that were running actual backup controller software (not the backup agents) were identified as 'Backup'

- A server directly involved in voice communications management was identified as ‘Communication’ (there were several VoIP PBX’s and voice-mail servers in the inventories)
- Anything that was running a named database server instance, regardless of the type of database, was identified as ‘Database’. Oracle, DB2, Sybase, ADABAS, and SQL Server (by far the most frequent) were identified in the inventories.
- All DHCP, LDAP, WINS, Active Directory, and DNS servers were categorized as ‘Directory’ servers (we understand that they serve other functions, but providing look-up services for names, access rights, addresses, etc. is one of the primary purposes of these servers).
- Anything that was running an e-mail services be it Exchange, SMTP Mail, Notes, etc., was identified as ‘E-mail’
- Any server providing either file transfer protocol (FTP) or network-attached storage and/or printing services was identified as ‘File/Print’.
- Servers that provided gateway services to legacy communications environments, such as SNA, SAA, etc., or proxy servers used for securing Internet access were identified as ‘Gateway’.
- Servers that were used to feed standardized executable image files for purposes of software installation were identified as ‘Image’
- Any server specifically used for software license management was identified as ‘License’.
- Anything that was used to monitor and manage other servers or networks was identified as ‘Management’.
- All Citrix and Microsoft Terminal Servers were categorized as ‘Remote Access’ servers.
- Anything that had to do with security (IP address checking, encryption, firewalls, etc.) was identified as ‘Security’
- Any server that could not be associated on some basis with any of the available categories was identified as ‘Unknown’.
- All servers running VMware or Microsoft Virtual Server were identified as ‘Virtual Hosts’ (virtual machines running on these servers were inventoried according to their individual primary function, separately, but identified as ‘virtual’ on the virtual/physical attribute).
- Any server providing virus scanning services was identified as ‘Virus’.
- Anything that was providing public or private Internet presence that was not specifically identified as an application web server was identified as ‘Web’

We did not review our categorization decisions with the providers of the original inventories due to schedule and scope constraints.

Ecora scan-based inventories provided comprehensive detail concerning the physical attributes of scanned servers and the software installed on them. However, the Ecora tool has no way of knowing whether or not a server is in production or development status, whether or not it is a line of business (LOB) vs. an infrastructure server, or what it’s primary function might be.

Since these attributes are important for analysis, we attempted to follow a two-phased process for the Ecora-based inventories:

1. Using the server name itself (the DNS host name from the Ecora scan); we attempted to sense whether or not the server was production or development, line of business or infrastructure. If the answer was obvious ('Dev' or 'Prod' were often found in the server name), then the attribute was marked accordingly. If not, we left it as 'Unknown'.
2. We manually scanned the list of installed software on a large number of Ecora-scanned servers. If we were able to identify a specific (non-utility) application, then we attempted to discover what that application did (usually via a Web search) and then used our judgment to assign a primary function to the server in question (and we marked the application instance as the 'primary' application in the Server_Application table). Otherwise, we set the status to 'Unknown'.

We did not attempt to review or reconcile the choices made as a result of this analysis with the site owners. We also were not able to finish the detailed review of all of the auto-scanned servers; the task proved to be too time-consuming.

Auto-scanned inventories are therefore the primary source of 'unknown' status for a server's primary function, category, and type within the database.

We made little attempt to standardize the server category (development, production, etc.) and server type (infrastructure, line of business) attributes. In general, we accepted whatever value the inventory respondents provided. Our 'sense of the data' is that while 'development' vs. 'production' was well understood, 'line of business' vs. 'infrastructure' was not. Therefore no conclusions should be drawn from this attribute. Use server primary function instead.

11.3. Appendix C: Physical and Virtual File Servers by Site

Page removed from report for security reasons.

12. Glossary of Terms

Term	Definition
Agency	A cabinet-level government entity. Consists of multiple departments.
Application servers	Servers which support particular applications, both custom and off-the-shelf. The distinction between a Web server and an Application Server, at least today, is becoming increasingly blurred, since most applications today are delivered via the Web. For this study, application servers were considered to be those that performed specific, identifiable applications, not just general Web service.
Data Center	Any location that is used specifically to contain multiple server computers and supporting infrastructure. All 'raised floor' computer rooms are data centers, but not all data centers incorporate raised floor and other sophisticated facilities.
Database	A server running a dedicated commercial database management system, such as Oracle, SQL Server, DB2, etc.
Department	The organizational level within the State of California government at which most IT assets are deployed and managed.
E-mail server	A server running a dedicated e-mail system such as Exchange or SMTP Mail.
Enterprise Content Management	A term that describes the emerging discipline of categorizing, storing and retrieving some/or all of the electronic content (as opposed to physical media, which goes by the moniker of 'records management') in an organization.
File/Print Server	A server used to provide a common storage location for many users and/or other servers for data files of various kinds. Also used to manage the print queues that are associated with printers. (Includes file-transfer protocol (FTP) servers.)
Frequency	Precisely: the number of times a processor's clock oscillates per second. Practically: a reasonably useful guide to relative performance between microprocessors of the same architectural design.
FTE	Full-Time Equivalent - FTE support includes all support activities including first level technicians, first, second, third level administration and technical support, troubleshooting, supervision and management, end user technical assistance.
GbE	Gigabit (one billion binary digits) per second Ethernet (a standard mechanism for physical transmission of data over a particular collection of wires or optical links.
Hyperthreading	A capability of certain Intel microprocessors to 'appear' to the operating system as if there are multiple physical processors present when in fact only one processor is present. Not to be confused with 'multicore', which does in fact provide more than one physical processor in a single socket. Hyperthreading must be enabled by the system BIOS in order to be used.

IDC	IDC is a global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets.
LAN	Local Area Network
Primary Function	For purposes of this study, the 'primary function' of a server is intended to capture the main purpose to which a particular server is dedicated. We recognize that many servers perform multiple functions, but for purposes of this study only one primary function was identified and captured for each server.
Proxy	2 Definitions: 1) a server which acts 'on behalf of' another server, typically for purposes of internet access (we categorized this as a 'gateway' server for primary function purposes and 2) equivalent to 'approximation', as in "the project team used the number of Windows and Intel-based dual-socket servers as the proxy for [the total number of] servers".
Raised Floor	In purpose-built computer rooms within data centers, raised floor is installed to simplify cabling and improve airflow throughout the facilities. Raised floor is more expensive to construct and maintain than conventional flooring.
SAN	Storage Area Network
Server	A general-purpose computer (as opposed to a dedicated 'appliance') which supports many simultaneous users.
TCO	Total Cost of Ownership. In this study, includes both acquisition costs and all costs associated with owning and operating an IT asset over the useful life of that asset.
Tier III	A designation (by a sponsoring professional organization) of a Data Center design as being continuously operable during both the planned and unplanned outage of any single underlying component.
U.S. Volume Servers	IDC term - Volume server market (consisting of all systems with an average sales value [ASV] below $25,000)
Virtual Machine	An instance of a server running on a physical machine in software. The physical machine can host one or more virtual machines.
Virtualization	In computing, virtualization means to create a virtual version of a device or resource, such as a server, storage device, network or even an operating system. For our usage, it is a representation of a real machine using software that provides an operating environment which can run or host a guest operating system. Multiple virtual machines, i.e. servers, can be run in software on one physical machine.
Web server	A server that responds to Internet- or Intranet-generated requests for Web pages. There are many different variations of 'Web servers' in use, depending on the specifics of a particular Web implementation architecture. All servers identified as having something to do with hosting Web pages were identified as 'Web servers'.

Wide-Area File System	A type of file service that explicitly takes into account the unique challenges of long-distance, relatively lower-bandwidth, higher-latency network connections that are typically of wide-area networks. WAFS are newer products which use intelligent caching, compression, and other techniques to improve performance of file systems even when they have been centralized over the WAN.
X86	The term generally applied to any microprocessor that is compatible with the instruction set originally defined by the Intel 8086 processor, circa 1981.